

GWAVA Inc.

GWAVA[®] Retain[™]

For GroupWise[™]

Administration and User Guide

Version 1.8

"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.

Contents

<i>Copyright Notices</i>	6
GroupWise Error Handling.....	7
<i>Initial Setup Overview</i>	8
<i>Running GWAVA Retain for the first time</i>	9
<i>About Retain Storage Paths</i>	10
<i>Advanced Storage Options</i>	10
<i>Retain – Initial setup</i>	11
<i>IMPORTANT Information about the Trusted Application Key generator</i>	12
Using GWAVA Retain.....	17
Logging In.....	17
<i>Archiving/Data Collection – About Workers, Profiles, Schedules and Jobs</i>	19
<i>Retain Job components:</i>	19
<i>Saving Changes</i>	20
Schedules.....	21
Profiles.....	22
<i>About GroupWise’s Smart Purge and Retention Flags</i>	26
<i>Setting up SmartPurge and Retention flags in GroupWise 7</i>	28
Workers.....	32
Jobs.....	39
Management.....	43
Groups.....	43
<i>Creating a New Group</i>	43
<i>Group Rights</i>	44
Users.....	47
<i>Offline Password</i>	47
<i>Creating a new user</i>	47
<i>Core Settings</i>	48
<i>User Rights</i>	49
<i>User-Level Rights Summary</i>	50
<i>The Read Configuration right (Redline Integration)</i>	50
<i>Mailboxes</i>	51
Address Book Selector.....	52
<i>Adding Mailboxes</i>	52
<i>About “Show only recently cached items”</i>	52
<i>Miscellaneous Tab</i>	53
<i>GroupWise Proxy Support</i>	54
Deletion Management.....	55
<i>Core Settings</i>	55
<i>Basic Options</i>	55
<i>Data Scope</i>	56
<i>Job Members</i>	57
<i>Notification</i>	57
<i>Schedule</i>	57

Configuration	58
Server Configuration	58
<i>Communications</i>	58
<i>Index</i>	59
<i>Storage</i>	59
<i>Accounts</i>	60
<i>Open System vs. Closed System</i>	60
<i>LDAP and E-Mail Address Based Login</i>	60
<i>Maintenance</i>	62
<i>Post Office Redirection Tables</i>	63
<i>Ignore Domains</i>	63
<i>Notification</i>	64
<i>Logging</i>	65
<i>Auditing Records</i>	65
<i>Miscellaneous</i>	66
<i>Export XML</i>	66
Storage Manager / Storage Engine	67
Partitions	70
Stubbing Server	71
<i>ConsoleOne GroupWise settings</i>	71
<i>Retain Stubbing Settings</i>	73
<i>Restore Stub functionality</i>	76
Licensing.....	77
Overview	78
System Log.....	78
Audit Log.....	79
Searching the Retain Archives.....	80
<i>Quick View</i>	81
<i>NOTE: Date Range Filter</i>	81
<i>Using Online Help</i>	81
Browse	82
<i>Language Selection</i>	82
<i>Browse Window Defaults</i>	83
<i>New Mailbox</i>	83
<i>Toolbar Options:</i>	84
<i>Specifying a Date Range</i>	84
<i>Forwarding a Message</i>	84
<i>PDF Export</i>	85
Search	88
Opening/Viewing a Message.....	89
Search in More Detail.....	91
<i>Search Window</i>	91
<i>Core Tab:</i>	91
<i>Selecting Mailboxes</i>	92

<i>Loading, Saving and Deleting Searches</i>	92
<i>Scope Tab:</i>	93
Item Type	93
<i>Sort Tab:</i>	93
<i>Misc. Tab:</i>	93
Search Window	94
<i>Core Settings</i>	94
<i>User Rights Tab</i>	95
<i>Mailboxes Tab</i>	95
<i>Miscellaneous Tab</i>	95
Using the Retain Help System	96
Credits and Diagnostics Page	97
Appendix A:	98
Backing up Retain	98
Fine tuning GWAVA Retain	99
<i>Tomcat Memory Usage Tuning Overview</i>	99
<i>Two Critical Issues</i>	99
General Recommendations	99
<i>Retain Server Component</i>	99
<i>Retain Worker Component(s)</i>	99
Tomcat Configuration Parameters	100
How to change memory values	101
<i>Checking memory:</i>	101
UPGRADED SYSTEMS	102
Appendix B – Troubleshooting Some Common Error Messages	103
<i>EA04</i>	103
<i>D712/D714</i>	103
<i>C05D</i>	103
<i>Other notes regarding GroupWise 7, SP 3 (7.0.3):</i>	103
Appendix C – The Tools Folder	104
<i>The Archive Migration Tool</i>	104
<i>The Archive Plug-in</i>	104
<i>Blob Extraction Tool</i>	104
<i>Diagnostic Tester</i>	104
<i>Recover Admin</i>	104
<i>Single Sign On</i>	105
<i>WebAccess</i>	106
For a text-only or shell install on Linux:	106
<i>System Requirements</i>	107
<i>Overview; Modes of Operation</i>	107
<i>How it Works</i>	107
<i>Prerequisites</i>	108
<i>Setting It All Up</i>	108
<i>Running the Migration</i>	111

<i>Limit of 5 Concurrent Migrations</i>	111
<i>Other Notes</i>	111
Appendix E – Reload Integration	112
<i>A Brief Review on How Reload Works</i>	113
<i>How Retain Takes Advantage of Reload's features</i>	114
<i>Multiple Retain Workers on One Server</i>	114
<i>Timing</i>	114
Retain Settings.....	115
<i>How to Setup Reload to work with Retain</i>	116
IMPORTANT Notes for the Integration	123
<i>Retain</i>	123
<i>Reload</i>	124
Appendix F – Migrating Retain to a new Server or Platform.....	125

IMPORTANT NOTES

Current Retain Customers, Retain 1.8 Upgrade

Retain 1.8 contains several important changes from Retain 1.5, and performs these updates the first time 1.8 is run. To run the update, run Retain 1.8 after the update and login as administrator. Retain will then perform the update. **DO NOT INTERRUPT THE UPDATE OR DATABASE CORRUPTION MAY RESULT.** Wait until the Retain console notifies you that the update has completed, then restart Tomcat.

Updates

Deletion Manager

Retain 1.6 and above include a new Deletion Manager option. This option will not be visible unless the administrator account is granted the Deletion Manager rights in user management.

There are two major updates to the system:

➤ **Indexer updated**

Lucene, the indexing engine that Retain uses to catalog and index the mail archive, has been updated from version 2.2 to 2.3. The update changes the index schema, providing performance enhancements. Because of the update, once 1.8 has been installed and run, it is impossible to return to a previous version, as the previous version will view the index as broken. **The upgrade is permanent.**

➤ **Database update**

There is a database update between 1.8 and previous versions. With large databases, the update can take hours. When you first connect to the Retain Server console after the upgrade, you will be prompted to initiate the database update. This is the only screen available until the database update is complete. **DO NOT INTERRUPT THE UPDATE OR DATABASE CORRUPTION MAY RESULT.** If any errors occur, Retain lists them in the Server logs. The Retain Server logs are nested with the embedded Tomcat 5 logs, which differs depending on installation platform.

Linux: /opt/beginfinite/retain/tomcat/logs

Windows: C:\Program Files\Apache Software Foundation\Tomcat 5.5\logs

Retain Server moving to 64-bit architecture

Retain Server, to accommodate the needs of large systems and stability, is moving to 64-bit architecture. 32-bit Server support will end with the Retain 1.x tree and **Retain 2.x will only support 64-bit server operation.** It is suggested that administrators begin planning their systems accordingly. The Retain Worker will still be supported on both 32-bit and 64-bit systems.

The list of supported operating systems is as follows:

SuSE Enterprise Linux 9 SP4 32-bit

SuSE Enterprise Linux 10 SP2 32-bit

Open Enterprise Server 1, 2 (Linux) 32-bit versions

SuSE Enterprise Linux 10 SP2 64-bit

SuSE Enterprise Linux 11 64-bit, (32-bit for worker only)

Open Enterprise Server 2 64-bit

NetWare 6.5 (Retain Worker only)

Microsoft Windows Server 2000, 2003 32-bit

Microsoft Windows Server 2003, 64-bit

Microsoft Windows 2008 64-bit

New Supported Software or features:

Retain also now supports 64-bit JVM and Centera for SLES 10 SP2, OES2, and Windows Server 2003.

Retain also supports NetApp.

Retain now performs LDAP email address lookups. For large eDirectory trees, email addresses will need to be indexed for this feature to work. GroupWise does not index these addresses by default.

Retain also supports auto-creation of users in Livelink.

Retain supports the Hitachi Content Archive Platform, (HCAP).

Retain Supports Caringo CASTor storage system

GroupWise 8 Support

Retain 1.8 provides full support for GroupWise 8

Java Support

The supported versions of Java have been unified for the Retain Server platform. **SUN Java 1.5 is the only supported version and is required for full functionality, (also called Java 5 on Windows) IBM's Java platforms are not supported.** Retain 1.7 will not work with Java 1.4x. It is strongly recommended that the JDK package is installed and used instead of the JRE package, for the diagnostic tools included with the JDK software package. Retain will automatically install Sun JDK 1.5 for you if the Retain Server has internet access.

Stubbing Server

Retain now has an integrated server to take advantage of the stubbing abilities of GroupWise 8. The Stubbing feature in GroupWise 8 allows Retain to replace archived messages and items in the GroupWise Post office that exceed a certain size with smaller 'stubs' that link back to the full message or item in the Retain Archive. The end user experience is not changed. Stubbing is not appropriate for every system. **Once enabled, the stubbing process cannot be reversed. To use stubbing, it is strongly recommended that the GroupWise system be at least 8.0SP1. Previous versions of GroupWise do not contain critical stability fixes.**

Archive Mark

Retain now can create a custom field to help users know which messages have been archived. This feature is not to be used as a backup compliance check, as it can be easily modified, but is simply for informational purposes. See the [Archive Mark](#) section under Profiles for information.

Copyright Notices

The content of this manual is for informational use only and may change without notice. GWAVA Inc. assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation. GroupWise is a registered trademark of Novell, and copyrighted by Novell. Windows is copyrighted by Microsoft. © 2010 GWAVA Inc. All rights reserved. ® GWAVA is a registered trademark. All other trademarks, service marks, and company names are properties of their respective owners.

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

GroupWise Error Handling

Retain 1.8 has better GroupWise error handling for errors on messages. These errors are encountered when Retain requests a message and any attachments, and the data is either corrupted or missing, causing the GroupWise Post Office to return a terminal error to Retain for that message part. These messages cause the same errors when accessed from the GroupWise client.

Due to compliance design considerations in previous releases, these GroupWise errors may cause mailboxes with ten or more occurring errors to be aborted. Also the storage time stamp was not advanced because Retain was unable to archive all messages in the mailbox.

Retain now handles the four most prevalent terminal message errors differently. When GroupWise returns one of the four specified error codes instead of a requested message, Retain will write an exception error to a special log file in the logging directory, replace the missing information or email with a text file note stating which error occurred, and move on to the next message. Retain will also advance the storage time stamp as appropriate. The text file that Retain adds to the archive in the place of the corrupted or missing message part will be named the same as the original message file.

Retain takes these actions according to the options stated in the errors.properties file. The actions may be changed and options are explained in the file text. DO NOT MODIFY the errors.properties file unless the default actions are unacceptable to organization policy.

The file is located at:

...retain/RetainWorker/WEB-INF/classes/config/errores.properties

Adding error codes to the original four codes is not appropriate, as the actions and options available are not appropriate for other errors. The four codes are included because they are terminal for the messages in the GroupWise mailbox; unless a backup of the database exists prior to the corruption date, the data no longer exists and it is impossible to archive. Other error codes which do not fall into this category should not be added to the list.

Important note:

When the upgrade option is selected from the Retain install menu, for 1.8 and later, all logs, including the special log, and custom changes to the errors.properties file will be overwritten and reset.

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 no matter what platform it is running on.

To get Retain ready for regular use, some last one-time procedures need to be run.

- Retain Server needs to be configured.
- At least one Retain Worker needs to be configured.

Initial Setup Overview

We recommend using a web browser on a Windows workstation to run the initial setup of Retain, since Novell's API only allows you to generate a Trusted Application Key on Windows. The workstation must also be running Novell's client. You will not require Windows after the initial setup.

- 1) Choose your storage path.
- 2) Create an embedded database used to store all the configuration information and the GroupWise address book. Anything non-message store related is stored here. This step doesn't require any user intervention.
- 3) Retain prompts you to 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 embedded 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, so strictly speaking it is incorrect to refer to an "admin" user, but rather a user possessing some specific administrator level rights.
- 4) Retain prompts you for SMTP information for notification and forwarding.

DNS host names are generally recommended instead of IP Addresses in almost all cases because they are easily centrally managed, and hence fixed 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) 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 will create the tables and do all the rest.
- 6) Retain prompts you to create a Trusted Application Key, using a utility you download directly from your Retain server (there is a hyperlink provided on the setup page). This utility is written for Windows, and requires a workstation with the Novell client. You will only need this workstation once.
- 7) Retain prompts you to specify a default Post Office Agent. This POA is used for all address book caching, and is the point of the first try for ALL queries until the workers learn where their

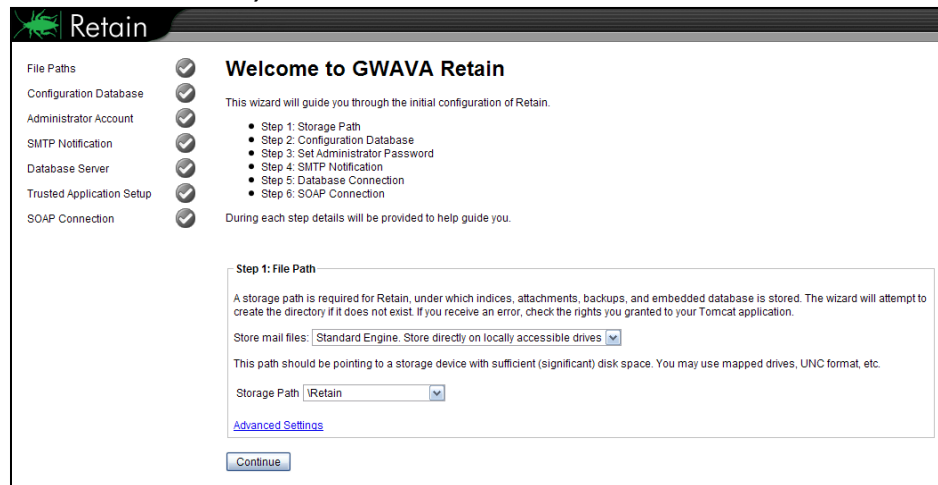
assigned PO is. It needs to be speedy and reliable!

- 8) Caches the GroupWise Address Book. Retain thus becomes aware of all GroupWise users in your system. This may take a few minutes, depending on the size of you GroupWise system, but typically takes only one or two minutes at most. The Address Book will be re-cached on a regular basis. See [Server Configuration](#).

Running GWAVA Retain for the first time

Note: Tomcat needs to compile and load Retain. On initial start or after an update allow Tomcat several minutes to load and run Retain; how much time depends on memory tuning and your system's setup. If you have precompiled your JSP's this startup is faster. See your platform's install guide for more information.

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



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

For descriptions of the different engines, see the [Storage Manager](#) section. If you select a storage system that is no 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".

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.

About Retain 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. Perhaps 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. Estimate your storage requirements. When choosing your storage system, bear in mind the following considerations:

- 2 TB limit. Will your 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. (Found after setup, 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.

Encryption

Compression

aes

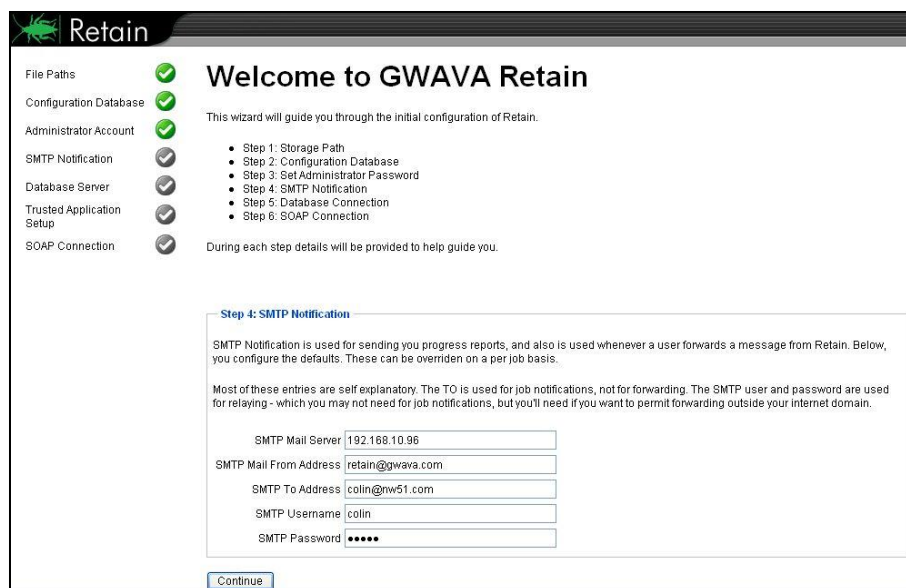
Retain – Initial setup

When you first call up the Retain Server screen, you are presented with this welcome page.



The screenshot shows the 'Welcome to GWAVA Retain' window. On the left is a sidebar with configuration steps: File Paths, Configuration Database, Administrator Account, SMTP Notification, Database Server, Trusted Application Setup, and SOAP Connection. The first three are marked with green checkmarks, and the last four with grey checkmarks. The main area is titled 'Welcome to GWAVA Retain' and contains a list of steps: Step 1: Storage Path, Step 2: Configuration Database, Step 3: Set Administrator Password, Step 4: SMTP Notification, Step 5: Database Connection, and Step 6: SOAP Connection. Below this, it says 'During each step details will be provided to help guide you.' The current step is 'Step 3: Administrator Account', which explains that an administrator account is needed for full access. It includes input fields for 'Login Name' (pre-filled with 'admin'), 'Password', and 'Confirm Password'. A 'Continue' button is at the bottom.

- 9) 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' window at Step 4: SMTP Notification. The sidebar on the left shows that 'SMTP Notification' is now the active step, marked with a green checkmark. The main area explains that SMTP Notification is used for progress reports and message forwarding. It includes a list of steps and a note that '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'll need if you want to permit forwarding outside your internet domain.' Below this, there are input fields for 'SMTP Mail Server' (192.168.10.96), 'SMTP Mail From Address' (retain@gwava.com), 'SMTP To Address' (colin@nw51.com), 'SMTP Username' (colin), and 'SMTP Password' (masked with dots). A 'Continue' button is at the bottom.

- 10) 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 MacAfee, Symantec, and SUSE).

Next, this screen appears for entering the database connection parameters. This is the same information you will have created when you made the database during the preparation phase. (See Install guide)

Step 5: Database Connection Parameters

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

- Driver: Type of Database Server. Currently only MySQL 5 is supported. MySQL 4.x will NOT work correctly.
- 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.)
- 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.
- 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 Username:

DB Password:

- 11) Next you will be asked to supply a trusted application key. If you haven't already done it, click on the link to download the application to your workstation and select RUN.

Retain

File Paths ☒ Configuration Database ☒ Administrator Account ☒ SMTP Notification ☒ Database Server ☒ Trusted Application Setup ☐ SOAP Connection ☐

Welcome to GWAVA Retain

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

- Step 1: Storage Path
- Step 2: Configuration Database
- Step 3: Set Administrator Password
- Step 4: SMTP Notification
- Step 5: Database Connection
- Step 6: SOAP Connection

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

Step 5a: 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.

- [Download the key generation application.](#)
- Run the key generation application
- Paste the key generated below

Trusted Application Key:

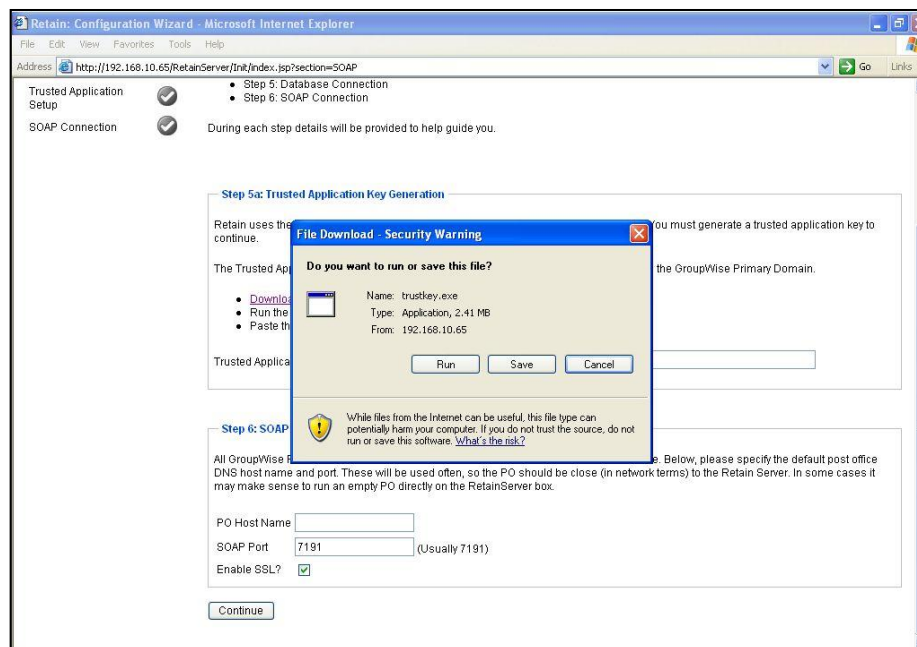
Step 6: SOAP Connection

All GroupWise Post Offices that Retain accesses must have SOAP enabled in ConsoleOne. Below, please specify the default post office

IMPORTANT Information about the Trusted Application Key generator

- 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.

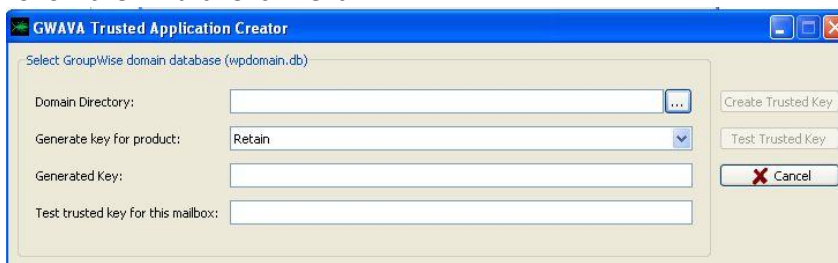
Note: The Trusted Application Key takes a few moments to propagate through the GroupWise system and therefore it may not be ready for testing or use immediately after you've generated it.



12) Click RUN to download and run the Trusted Application Key generator.



13) Follow the wizard. Click Next.



14) Find your GroupWise domain database. Click on the “...” button.

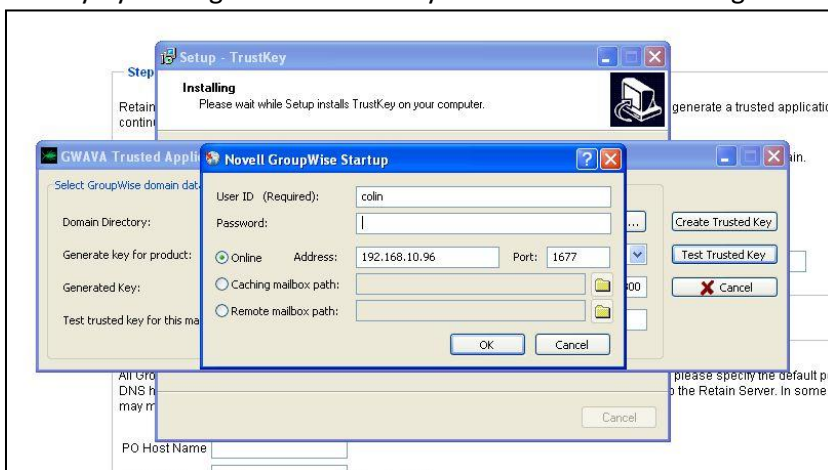


15) Once you have found your domain, click “Create Trusted Key”.



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.

16) 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.

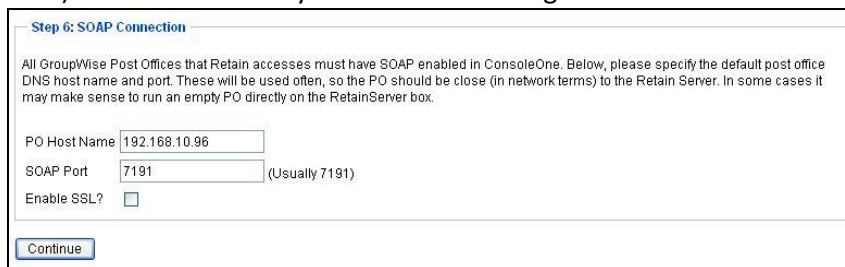




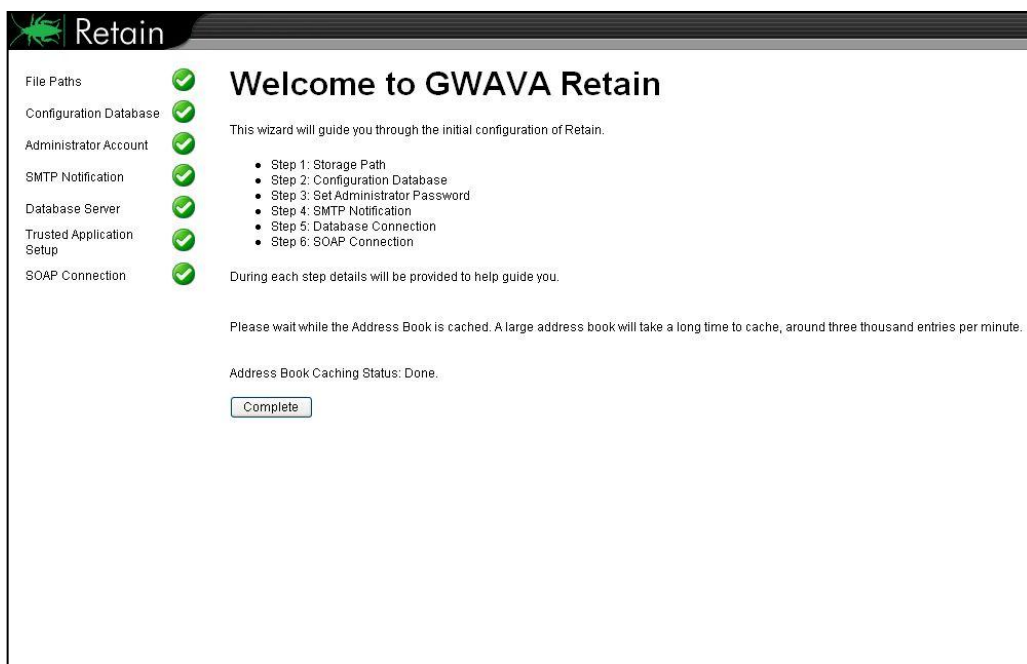
- 17) Once your login was successful, you know that both your trusted application key and the SOAP mechanism are working properly. The trusted application key is now copied to your clipboard so you can paste it to the Retain configuration page.



- 18) Now, below the Trusted Application Key section, enter the Post Office address (hostname/DNS or IP) and double-check your SOAP SSL setting. Click "Continue"



NOTE: Using a DNS name instead of an IP address is strongly recommended.




- 19) When all initial setup steps have been completed, a “Complete” button will appear. Click it and will be returned to a login page. You are now ready to move on to the next section, [Using Retain](#) .

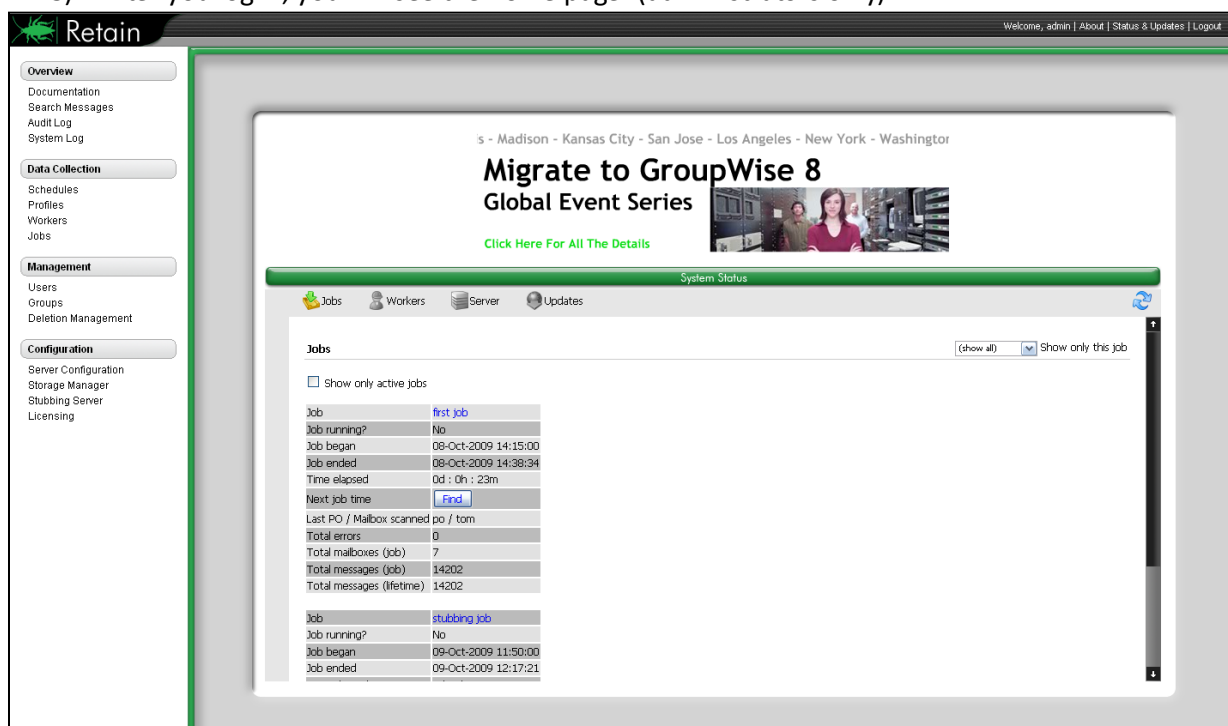
Using GWAVA Retain

Logging In

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



- 3) After you log in, you will see the home page: (administrators only)



This is the “status and updates” page. The latest status of jobs and recent activity is displayed on this page. If updates to the software are available, it’s displayed here too.

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).

The status page has four different tabs which show the status of your system. The status page is set to show what is deemed the most important for information at any given time.

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 on 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.

The screenshot shows the 'Server' tab selected in the top navigation bar. The main content area displays the following information:

SQL Server:	OK
Server:	RetainServer (ECMBDJORFRUJXBTJEFODSH)
Stored EMails:	0
Server Errors:	0
Installation Time:	12-Aug-2008 13:58:37
Enable Address Book Caching:	Yes
Last Address Book Cache:	27-Aug-2008 01:02:10
Enable Backup of Embedded Database:	Yes
Last Backup of DB:	17-Aug-2008 01:02:06
Enable Compression of Embedded Database:	Yes
Last Compression:	27-Aug-2008 01:02:01
Enable Backup of Indices:	Yes
Last Backup of Indexes:	17-Aug-2008 01:02:07
Enable Index Optimization:	Yes
Last Index Optimization:	27-Aug-2008 01:02:02

The screenshot shows the 'Workers' tab selected in the top navigation bar. The main content area displays the following information:

Worker:	linux worker
Last Contact:	27-Aug-2008 21:23:45
Last Job:	job1
Last Job Start:	21-Aug-2008 08:23:28
Last Job End:	21-Aug-2008 08:23:49
Registered:	13-Aug-2008 16:38:20
Total messages stored:	60
Total mailboxes scanned:	12

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.

The screenshot shows the 'Jobs' tab selected in the top navigation bar. The main content area displays the following information:

Jobs (show all) Show only this job

☐ Show only active jobs

Job	new job
Job Running?	No
Job Began	15-Aug-2008 18:03:12
Job Ended:	15-Aug-2008 18:03:26
Time elapsed:	0d : 0h : 0m
Next Job Time	15-Aug-2008 13:10:00
Last PO/Mailbox Scanned:	po / admin
Total Errors:	11
Last Error:	
Total mailboxes (job):	4
Total messages (job):	20
Total messages (lifetime):	40

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.

The screenshot shows the 'Updates' tab selected in the top navigation bar. The main content area displays the following information:

Updates

You have the latest build.

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 archive policy will cover several key points:

- ✓ **What** to archive (mailboxes/jobs)
 - Includes what post office(s) to collect.
- ✓ **Who** will archive it (worker)
- ✓ **Where** will the data be stored (database configuration)
- ✓ **When** will the data be collected (schedule)
- ✓ **Why** you are collecting this data will govern **HOW** you collect it. It's about your goals.
 - Desired data retention policy
 - First time dredge or day-to-day update
 - What do you want to be able to search and retrieve?
- ✓ **How** to collect the data will cover (profile)
 - Types of items or users to archive.
 - Date range of items to be archived.
 - Avoiding publishing e-mails multiple times.
 - Using the Retention, SmartPurge, or Item Store 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. Simple options of a single time, (job runs once), or a recurring time frame are available.

Profile

The Profile tells the system which kinds of files to retrieve from the GroupWise system. All GroupWise 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 GroupWise SmartPurge and Retention flags before you implement them in your system.

Worker


The worker is the module of Retain that actually pulls the data from the Post Office. 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 Post Office, selected user list, or distribution list. The Worker will not start archiving anything until a Job is created which binds a Post Office 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.

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.

If you make changes inadvertently, or do not want to save the changes you have made, or wish to view the settings as they are on default in the system, then reload or browse off the page and back to the page you wish to view. This will reset the page to default.

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.

ONE-TIME 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.

The screenshot shows a window titled "Schedules" with a list of schedules on the left containing "default" and "new default". On the right, there are buttons for "Add Schedule", "Edit Schedule", and "Remove Schedule". Below the list, the "Schedule Name" is set to "default" and the "Schedule Frequency" is set to "Single". A message states: "No jobs are associated with this schedule." Below this, a note says: "This job will run only once, at the time specified below, unless an error occurs. If an occur occurs, the job will be retried unless you select the checkbox below." The date is set to "Aug 13 2008" and the time is "16:36". At the bottom, there is a checkbox labeled "If time assigned for job start has elapsed, don't resend job" which is currently unchecked.

RECURRING schedules are configured using a start and interrupt time, and which days to start on.

The screenshot shows a window titled "Schedules" with a list of schedules on the left containing "beta5". On the right, there are buttons for "Add Schedule", "Edit Schedule", and "Remove Schedule". Below the list, the "Schedule Name" is set to "beta5" and the "Schedule Frequency" is set to "Recurring". Below this, there are checkboxes for days of the week: Sunday (unchecked), Monday (checked), Tuesday (checked), Wednesday (checked), Thursday (checked), Friday (checked), and Saturday (unchecked). To the right of these checkboxes, there are fields for "Start Job at:" (12:00), "Maximum job duration" (22 hours), and a checkbox for "Do not interrupt job, even when duration expires" which is checked.

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 it what to archive.

The profile defaults to collecting all data, with no date restrictions. This is almost never a wise selection in practice. After you've primed Retain by collecting data the first time, you'll want to narrow the date range and criteria considerably, or pay an intolerable price in performance.

The screenshot shows the 'Profiles' configuration window. At the top, there's a search bar and a list of profiles. The 'first profile' is selected, showing 'stubbing enabled'. To the right of the list are buttons for 'Add Profile', 'Edit Profile', and 'Remove Profile'. Below the list, a message states 'Profile first profile loaded.' The 'Profile Name' field contains 'first profile'. Below this are tabs for 'Core Settings', 'Message Settings', 'Scope', 'Miscellaneous', and 'Advanced'. The 'Core Settings' tab is active, showing 'Profile Functions'. The text explains that Retain normally archives mail, but GroupWise 8 introduces 'Stub' items. It advises enabling Archiving and Stubbing separately or together. It also notes that stubbing on GroupWise 7 post offices will be ignored. The settings include: 'Enable Archiving (store messages in Retain)' (checked), 'Enable Stubbing (if items are already archived)' (unchecked), 'Do not stub Calendar Items (Appointments, Notes, and Tasks)' (checked), and 'Stub only items exceeding a size threshold' (unchecked). A 'Size Threshold (bytes)' field is set to '2048'. Below this is the 'Archive Mark' section, which includes the text 'You can request Retain to mark items as archived. You may view the status at the Windows client as a custom column. This does slightly degrade performance.' and a checkbox for 'Mark emails as archived?' (unchecked).

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 is a new feature, (requires GroupWise 8), which 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 8.0HP1. 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’.
- User’s 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.)
- Access to stubbed items is currently only available to the Windows Client.
- Non-mail items should not be stubbed. (Calendar and appointment items behave erratically when rescheduled).
- Stub items are NOT COMPATIBLE with third party API’s. These applications view the stubs as posted items only, with no text or attachments.



Determine the correct option for your system. If you wish to enable Stubbing, you need to install and configure the Stubbing Server as well. Refer to the install guide for your specific OS, and the [Stubbing Server](#) 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' Hoggins	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.

The screenshot shows the 'Message Settings' tab selected among five tabs: Core Settings, Message Settings, Scope, Miscellaneous, and Advanced. The settings are organized into four sections:

- Mailbox Type:** A text box explains that items can be stored from Users or Resources. Below are two checked checkboxes: ☒ Users and ☒ Resources.
- Item Type:** A text box states that by default, Retain stores every type of message. Below are five checked checkboxes: ☒ Mail, ☒ Appointment, ☒ Note, ☒ Task, and ☒ Phone Message.
- Item Source:** A text box explains that messages are classified into categories. Below are four checked checkboxes: ☒ Received, ☒ Sent, ☒ Draft, and ☒ Posted.
- Message Status:** A text box explains that storage can be restricted based on GroupWise status flags. Below are three dropdown menus: 'Opened is' set to 'Doesn't matter', 'Read is' set to 'Doesn't matter', and 'Private is' set to 'Doesn't matter'.

By default, all items are selected for collection.

Scope

The screenshot shows the 'Scope' tab selected among five tabs: Core Settings, Message Settings, Scope, Miscellaneous, and Advanced. The settings are organized into two sections:

- Date Range to Scan:** A dropdown menu is set to 'All Messages (ignore date)'.
- Duplicate Check:** A text box explains that Retain stores a single instance of each message and attachment, and that defining how the Worker determines new items is a performance factor. Below is a dropdown menu set to 'Try to publish all messages (SLOW)'.


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




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, so there is already a check for duplicates at the storage end. However, sending megabytes of messages and attachments to the Retain Server only to have them discarded because “they are already here” wastes the Retain Worker’s 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:



“Relaod” 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.

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:

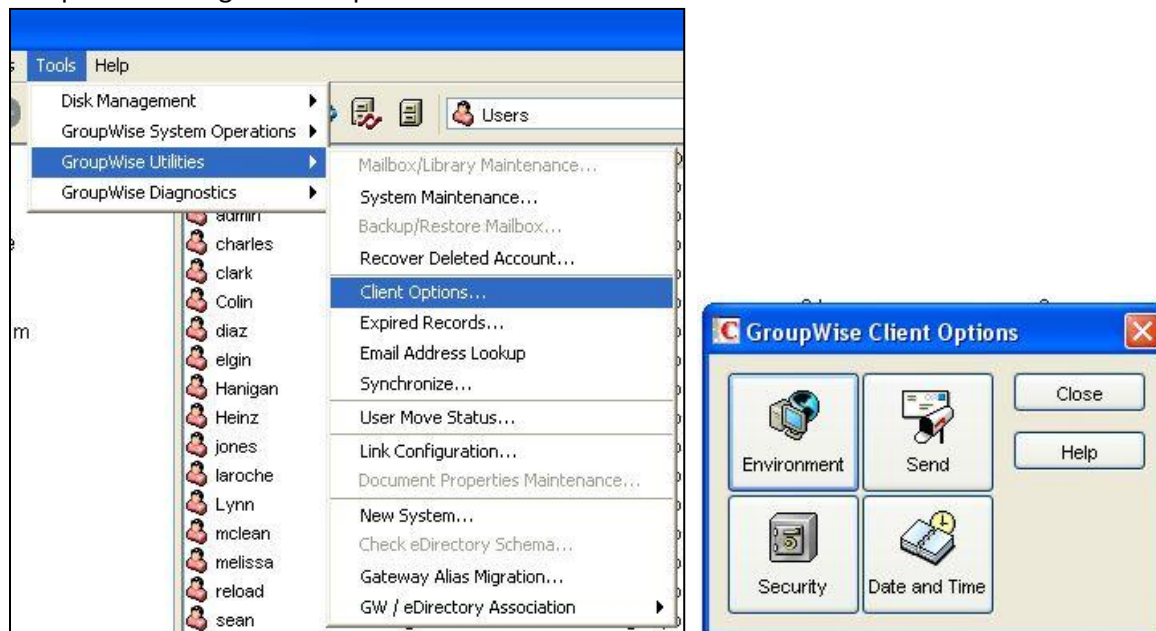
- The first is if for some reason the administrator needs to run multiple jobs or profiles on mailboxes that overlap. In this case, two separate flags are needed so the two different jobs can keep track of where (or more accurately WHEN) to continue archiving from. Admittedly, this is a rare scenario, 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

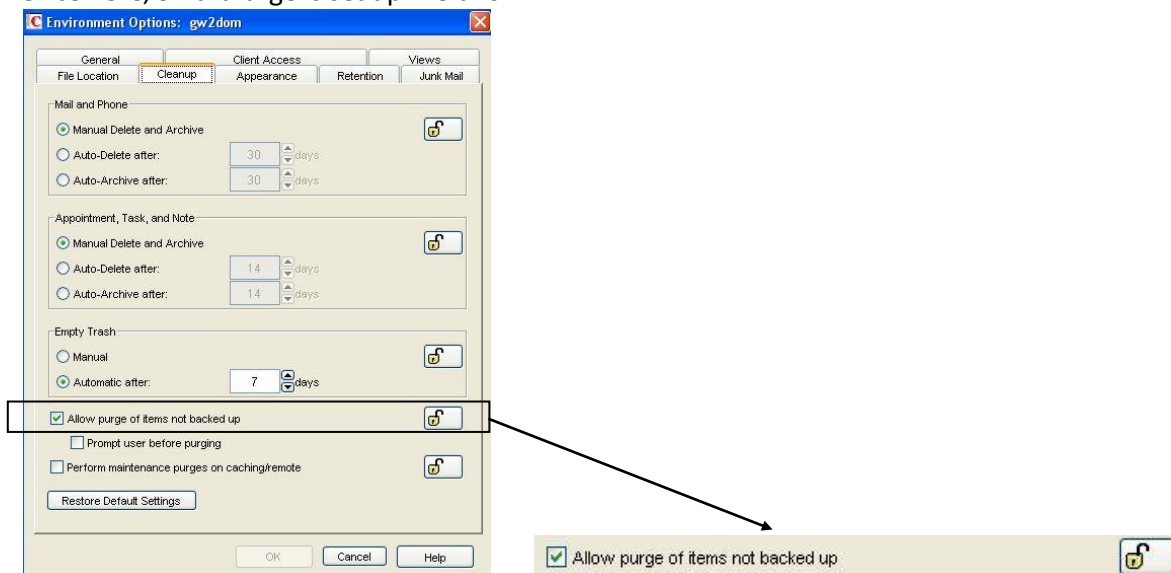
(Smart Purge Setup under GroupWise 8 is similar.)

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

Once here, SmartPurge is set up like this:

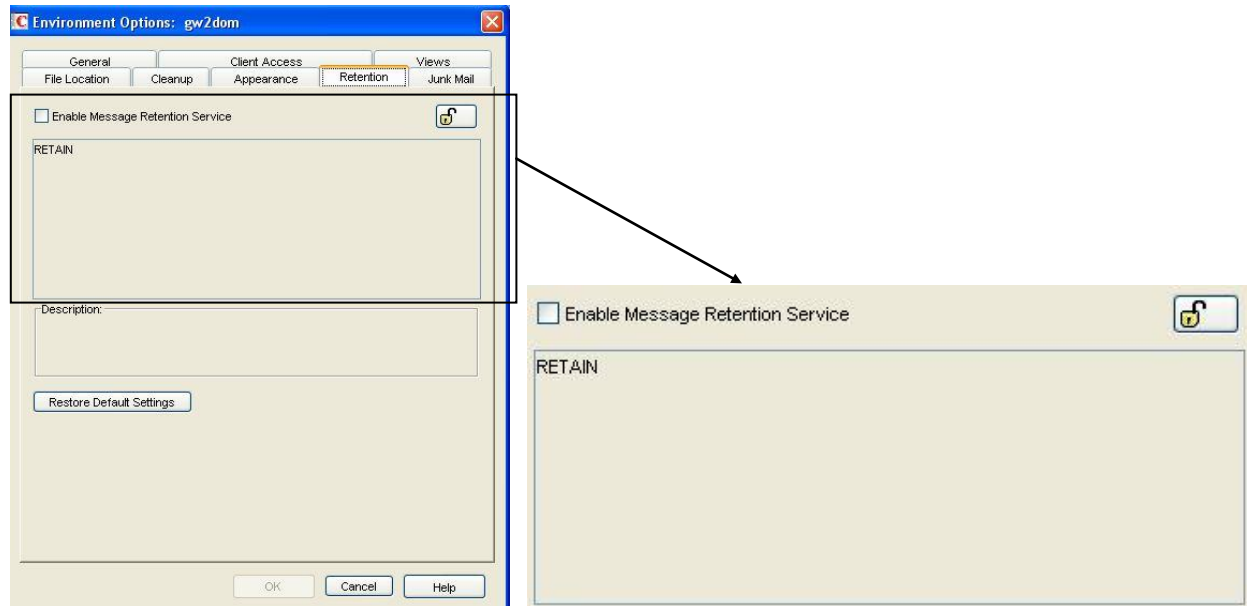


In the Cleanup tab, UNCHECK the box on “Allow purge of items not backed up” to ENABLE SmartPurge. Lock it to prevent users from disabling it.

(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.)

The Retention flag is set up a little differently than SmartPurge.

- 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.



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 TAB

The screenshot shows the 'Miscellaneous' tab in the GWAVA Retain configuration interface. At the top, there are five tabs: 'Core Settings', 'Message Settings', 'Scope', 'Miscellaneous' (which is selected), and 'Advanced'. Below the tabs, the 'Miscellaneous' section is titled in green. It contains two main areas. The first area, titled 'Miscellaneous' in green, explains that disk space can be saved by restricting attachment storage and that routing properties can be attached to messages. It features three radio button options: 'Store all attachments' (selected), 'Don't store MIME.822 attachments', and 'Don't store any attachments, other than the message'. The second area explains that users can choose to store Sent Items properties and/or store and index Internet Headers, noting that storing headers may have a performance penalty but is worthwhile for compliance. It includes two checkbox options: 'Store/index Internet Headers' and 'Include Routing Properties?'. The third area, titled 'Set Storage Flags' in green, explains that if using Purge or Retention features in GroupWise, these should be advanced automatically. It includes three checkbox options: 'Retention Flag' (checked), 'Purge Flag', and 'Item Store Flag'.

Core Settings Message Settings Scope Miscellaneous Advanced

Miscellaneous

You can save disk space by restricting attachment storage. You can also attach the Routing Properties to the message, if desired.

☒ Store all attachments
☐ Don't store MIME.822 attachments
☐ Don't store any attachments, other than the message

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).

☐ Store/index Internet Headers
☐ Include Routing Properties?

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

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.

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.*

ADVANCED TAB

The screenshot shows a web interface for configuring search criteria. The top section, titled "Advanced Criteria", contains an "Add" button and a "Delete" button. Below these is a dropdown menu currently showing "Subject", with a list of options: Subject, Sender, Recipient, Size, and Attachment Name. To the right of the dropdown is a text input field and a "is" dropdown menu. The bottom section, titled "Folder Scope", contains a dropdown menu currently showing "Items from All Folders", with a list of options: Items from All Folders, ONLY items from folders listed below, and All folders EXCEPT those listed below.

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 exdusions.

Your choices are:

- Dredge everything
- Dredge only these listed folders
- Dredge everything EXCEPT these listed folders

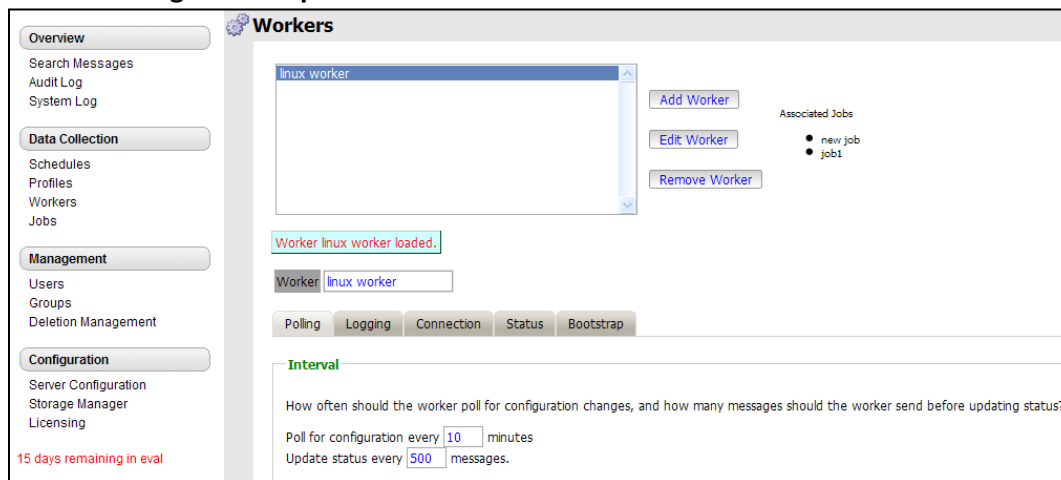
How to specify the list of folders to dredge/exdude:

- 1) Specify a System Folder (mandatory). Example: Calendar.
- 2) You specify a subfolder of that folder (optional).
Example: entering “old” would mean the folder “old” under “Calendar”.
- 3) You can have multiple hierarchies under that with the / delimiter.
Example “old/mail” would mean the subfolder “mail” under “old” under “Calendar”.
- 4) You specify if the option includes subfolder.
Example: If you select “old” and “includes subfolder” is unchecked, “Calendar/mail” is selected. If “includes subfolder” is CHECKED, “Calendar/old/mail” would also be selected.

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 Post Office Agent and delivering it to the Retain Server. The Retain Worker component is normally installed on each machine with a GroupWise post office 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 mailboxes) systems or for evaluation purposes. Ideally, the worker you choose to dredge mail from your post office is local to the post office itself. Using a worker on "Server A" to pull data from a post office on "Server B" is possible, but it will use up network bandwidth and it will be slower than simply pulling data from a post office on the same server. For this reason, Retain was designed to allow you as many workers as you need, on any platform supported by GroupWise.

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.
 - Protocol
 - URL/host name
 - password
 - If the connection information is changed, you will need to manually copy the RetainWorker.cfg bootstrap file over to the Worker
 - **Correcting the Bootstrap:**
 - 1) Stop Tomcat.
 - 2) First delete the RetainWorker.cfg file on the Retain 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: <retainworker>/WEB-INF/cfg
 - Follow the [Retain Worker Setup procedure](#).
 - GroupWise Soap Access

- If there is a reason to specify the connection address, (if the DNS does not correctly route to the GroupWise Post Office, or the internal address will not resolve to the Post Office), then you can manually set the connection address that the Worker will use to initially contact the GroupWise system. If this is left as default, the address used will be the initial SOAP connection you specified for the Retain Server to use. This setting is ignored if the Reload integration is selected.
- **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.

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

Workers

worker

Add Worker

Edit Worker

Remove Worker

No jobs are associated with this worker.

Enter new worker below

Worker new worker

Polling Logging Connection

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.

Worker Status	
Registered	12-Oct-2007 04:09:02
Last Contact	16-Oct-2007 22:50:46
Last Job	1576F3E6A27EDA9D0FF13C450C1070A7
Last Job Start	16-Oct-2007 15:11:45
Last Job End	16-Oct-2007 12:42:04
Total mailboxes scanned	60
Total messages stored	246496

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.

Server Protocol	http
Server Host Name	10.1.1.103
Server Port	80
Server Path	/RetainServer
Worker Password

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

GroupWise SOAP Access

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 Server Configuration, but it can be changed on a per worker basis.

POA Host Name	10.1.1.101
SOAP Port	7191
Enable SSL?	<input type="checkbox"/>

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.

GroupWise SOAP Access

If there is a reason to specify the connection address, (if the DNS does not correctly route to the GroupWise Post Office, or the internal address will not resolve to the Post Office), then you can manually set the connection address that the Worker will use to pull messages. This connection address will be the default connection to the POA that the worker will contact first regardless of which POA is assigned to be archived by the worker. If this is left as default, it will be the initial SOAP connection you specified for the Retain Server to use. *You should not need to change this setting for most systems.* For large or complex systems it can improve performance to set the SOAP Access address to the address of the POA which this particular Worker will be using. This setting is ignored if the Reload Integration setting is enabled.

Double check all these settings to ensure that you will have the best performance and connection for your system.

- 4) After you have checked and set your settings, click "Save Changes" before browsing to another page to create and save your worker.

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.

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

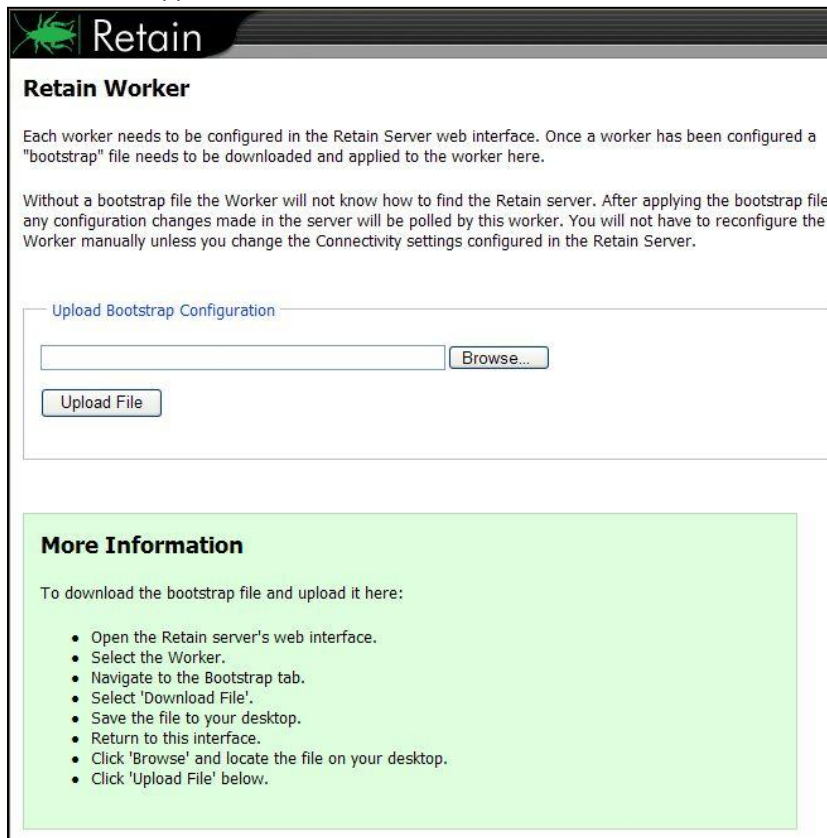


- 6) Save this file to your hard drive in a place where you can find it. Click “Save”.



- 7) Open a new browser window and type:
<http://<your-Retainworker-ip>/Retainworker>

This screen appears when a worker has not been initialized.



Retain Worker

Each worker needs to be configured in the Retain 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 Retain 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 Retain Server.

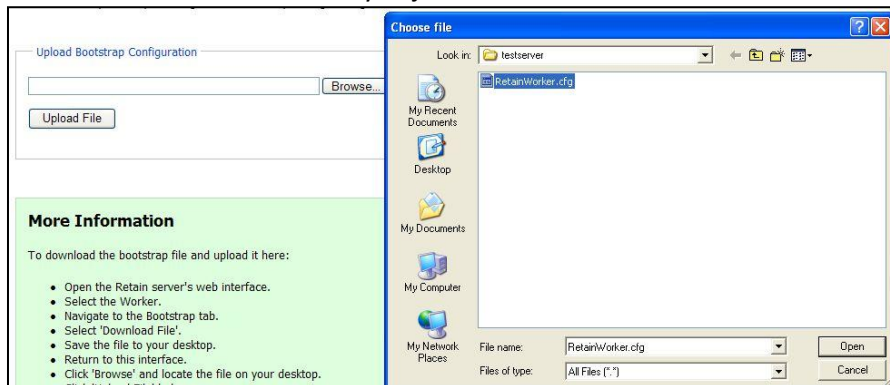
Upload Bootstrap Configuration

More Information

To download the bootstrap file and upload it here:

- Open the Retain 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.

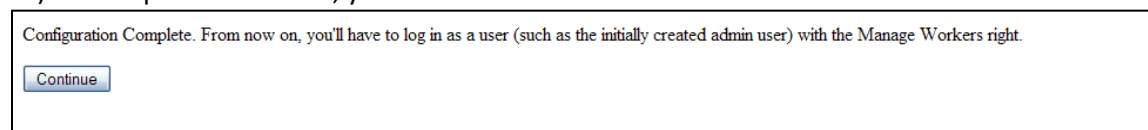
8) Click "Browse" to find the file you just downloaded to select it.



9) Click "Open".

10) Click "Upload File".


11) If the upload went well, you will see this screen:



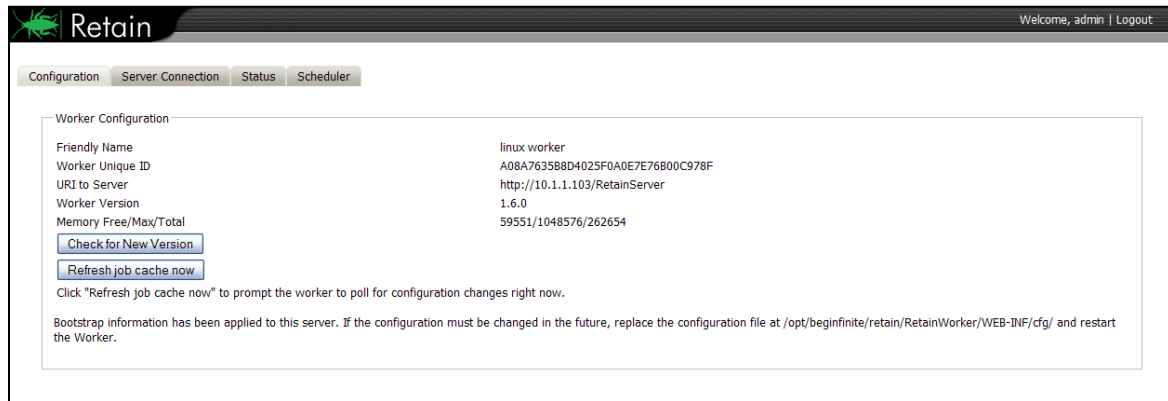
Configuration Complete. From now on, you'll have to log in as a user (such as the initially created admin user) with the Manage Workers right.

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

12) Login as Admin.




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



You should be able to see your configuration on the screen, as shown. The “Check for New Version” button checks for updates. The “Refresh job cache now” button tells the worker to poll the server for configuration changes and new jobs. After this button is selected, the Worker page will refresh automatically after the Retain Server has been polled. It takes just a few seconds.

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

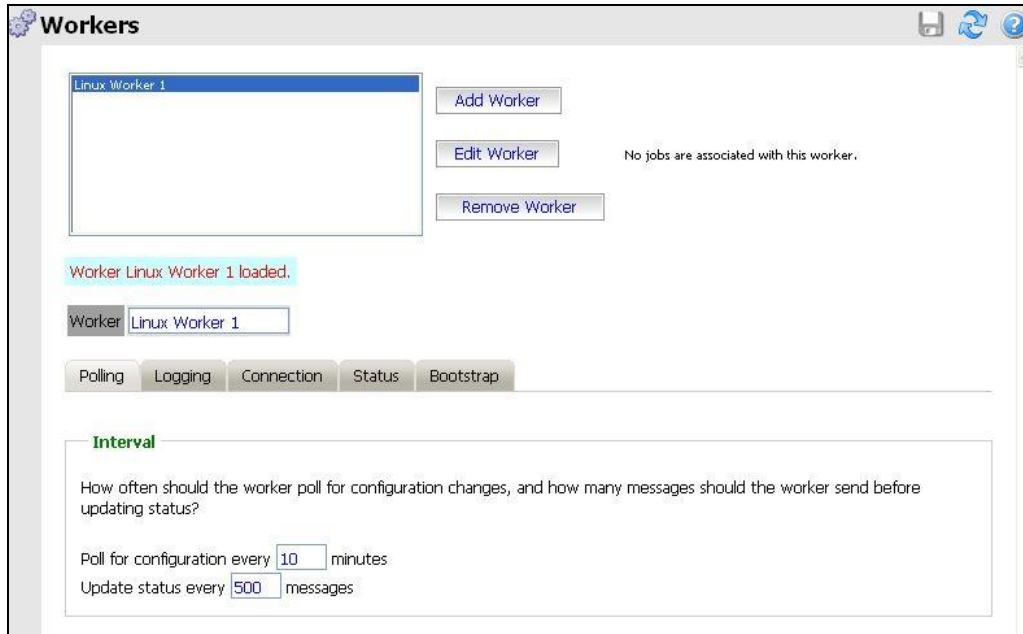


15) Go back to the Retain Server.

16) Click “Workers”.

17) Click on the worker you just created.

18) 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. Be aware of this when you submit a job. If the worker has not polled the server for the new job before the start time expires, the job will not start.

Also, status updates are displayed on the admin page every ‘n’ messages.

- 19) 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.
- 20) Now you need to set up [profiles](#), [schedules](#) and [jobs](#) to begin actual data collection.

Jobs

The actual data collection in Retain is controlled by **JOBS**.

- They start when their **SCHEDULE** says so.
- They collect data according to a specified **PROFILE**.
- They specify which **WORKER** will do the actual data collection.
- You specify which **POST OFFICE(s)** or **DISTRIBUTION LIST(s)** to collect the data from.

Jobs tie everything together. When you define a job, you specify:

- ✓ Schedule
- ✓ Profile
- ✓ Worker
- ✓ Entities to be scanned (Post Office(s), Distribution List(s))
- ✓ Notify in case of error
- ✓ Send Summary

Core Settings: Self Explanatory

The screenshot shows the 'Jobs' configuration window in GWAVA Retain. It includes a list of jobs, buttons for 'Add Job', 'Edit Job', and 'Remove Job', and a 'Core Settings' section. The 'Core Settings' section contains a 'Job enabled?' checkbox (checked), a list of required assignments (schedule, profile, worker), three dropdown menus for Schedule, Profile, and Worker (all set to '(none)'), a section for 'Enable Job Expiration?' with an unchecked checkbox, an 'Expire in' field set to '0' days, and a 'Base Expiration Time on' dropdown set to 'Date Item is Stored in Retain (recommended)'.

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, as well as provides support for NetApp SnapLock. You **MUST** enable and use job expiration for SnapLock to work.

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.

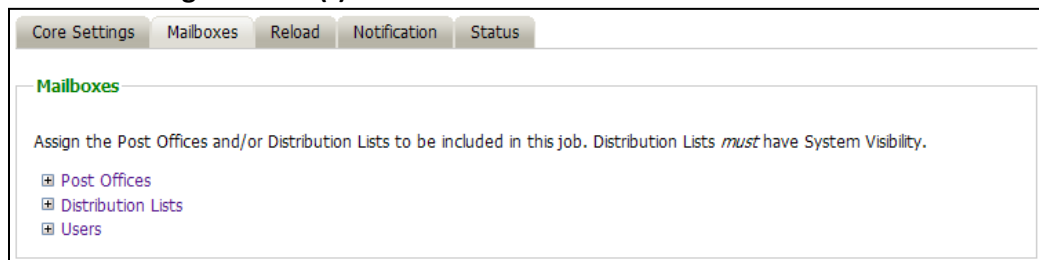
Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for.

Mailboxes

The mailboxes tab is where the administrator specifies which entities (Post Offices(s) and/or Distribution List(s)) are to be scanned.

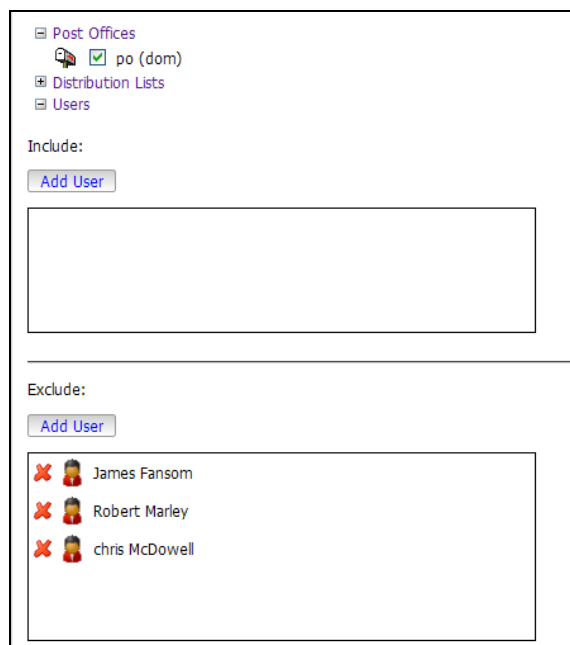
Expand the Post Office and/or Distribution List trees, and check off the items you want to be dredged.

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 screenshot shows the 'Mailboxes' tab in a web interface. At the top are tabs for 'Core Settings', 'Mailboxes', 'Reload', 'Notification', and 'Status'. Below the tabs, the 'Mailboxes' section is titled. A message states: 'Assign the Post Offices and/or Distribution Lists to be included in this job. Distribution Lists *must* have System Visibility.' Below this message are three expandable sections: 'Post Offices', 'Distribution Lists', and 'Users', each with a plus icon.

The users section allows you to select individual users to include, or exclude them from an archive job.



The screenshot shows the 'Users' section of the configuration page. On the left, there is a sidebar with a tree view containing 'Post Offices', 'Distribution Lists', and 'Users'. The 'Users' section is expanded. Below the sidebar, there are two main sections: 'Include:' and 'Exclude:'. Each section has an 'Add User' button. The 'Include:' section is currently empty. The 'Exclude:' section contains a list of three users, each with a red 'X' icon and a name: 'James Fansom', 'Robert Marley', and 'chris McDowell'.

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.

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' configuration page. At the top, there is a navigation bar with tabs: 'Core Settings', 'Mailboxes', 'Reload', 'Notification', and 'Status'. The 'Reload' tab is selected. Below the navigation bar, the page title 'Reload' is displayed in green. A descriptive text states: 'You can backup your GWAVA Reload server, saving network bandwidth if desired. You'll need Reload 2.5, and some settings set on the Reload machine. Please consult the user manual for additional information.' Below this text, there are two input fields: 'Enable Reload Integration?' with an unchecked checkbox, and 'Reload Server URL' with a text box containing 'http://myReloadBox:5555'.

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' configuration page. At the top, there is a navigation bar with tabs: 'Core Settings', 'Mailboxes', 'Reload', 'Notification', and 'Status'. The 'Notification' tab is selected. Below the navigation bar, the page title 'Notification' is displayed in green. A descriptive text states: 'You can have Retain e-mail a summary of operations and/or of errors that occurred when the job completes.' Below this text, there are several input fields: 'Mail when Errors Occur?' with an unchecked checkbox, 'Mail summary when job complete?' with an unchecked checkbox, 'SMTP To Address' with a text box containing 'admin@bitter.com', 'SMTP Mail From Address' with a text box containing 'Retain@linux.com', 'SMTP Mail Server' with a text box containing '10.1.1.101', 'SMTP Username' with a text box containing 'chris', and 'SMTP Password' with a text box containing four blue dots. At the bottom left, there is a 'Test Connection' button.

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	rpo	rpo	
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.

Management

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.

The screenshot shows the 'SPECIAL ADMIN PROPERTIES' interface for a group named 'auditor'. At the top, it states 'This group currently has administrative rights. This means:' followed by a bullet point: 'The members can see all audit logs.' Below this is a message 'Group auditor loaded.' and a text input field containing 'auditor'. There are four tabs: 'Core Settings', 'Group Rights', 'Mailboxes', and 'Miscellaneous'. The 'Group Membership' section shows a list with 'guest' and a 'Remove Member' button. Below the list is a 'Member' dropdown menu set to 'admin' and an 'Add Member' button. At the bottom, there is a 'Language' dropdown set to 'English' and a checkbox 'Prevent members from changing password?' which is checked.

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'. Under 'Administrator level rights', there are six unchecked checkboxes: 'Access all audit logs', 'Deletion Manager', 'Manage Server', 'Manage Users and Groups', 'Manage Workers, Schedules, Profiles, Jobs', and 'Search all mailboxes'. Under 'User level rights', there are five checkboxes: 'View/Save attachments' (checked), 'View personal audit log' (unchecked), 'Forward Messages' (checked), 'Print Messages' (checked), and 'Read Configuration (Redline)' (unchecked).

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 H4M2P1

Mailboxes

The screenshot shows the 'Mailboxes' configuration page for a group named 'default'. At the top, there are tabs for 'Core Settings', 'Group Rights', 'Mailboxes' (which is selected), and 'Miscellaneous'. Below the tabs, the 'Additional Mailboxes' section contains a text box with instructions: 'These are additional mailboxes beyond the primary mailbox that the members of the group 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 is a 'Currently Selected' section, which is currently empty. The 'Criteria' section includes a dropdown for 'First Name', a 'begins with' dropdown, a text input field, and a 'Search' button. Below the search criteria are two checkboxes: 'Show only users' (checked) and 'Show only recently cached items' (checked), followed by an 'Add Selected Items' button. The 'Address Book' section has a text box with the instruction 'Enter criteria to search the cached GroupWise 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 others' 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: 28-Aug-2008 19:48

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 3 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](#).

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 GroupWise users logging in with their GroupWise credentials
 - These users use SOAP authentication.
 - Retain checks their login credentials with GroupWise
 - 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 GroupWise
 - Users created in Retain do not need to have a GroupWise account.
 - Users who don't exist in GroupWise use the *offline password*.

Offline Password


There are, however, occasions when you might want someone to search through the Retain archives who is not part of the GroupWise 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 or via SOAP, the same rights can be assigned in either case. An administrator who possesses the Manage Users and Groups' administrative right can.

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 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.

The screenshot shows the 'Users' management window. At the top, a list of users is shown with 'admin' selected. To the right of the list are buttons for 'Add User', 'Edit User', and 'Remove User'. Below the list, a message says 'User admin loaded.' and a text field shows 'User ID: admin'. Below this are tabs for 'Core Settings', 'User Rights', 'Mailboxes', and 'Miscellaneous'. The 'Core Settings' tab is active, showing configuration options for the selected user.

The screenshot shows the 'Core Settings' configuration panel for a user. It includes the following fields and options:

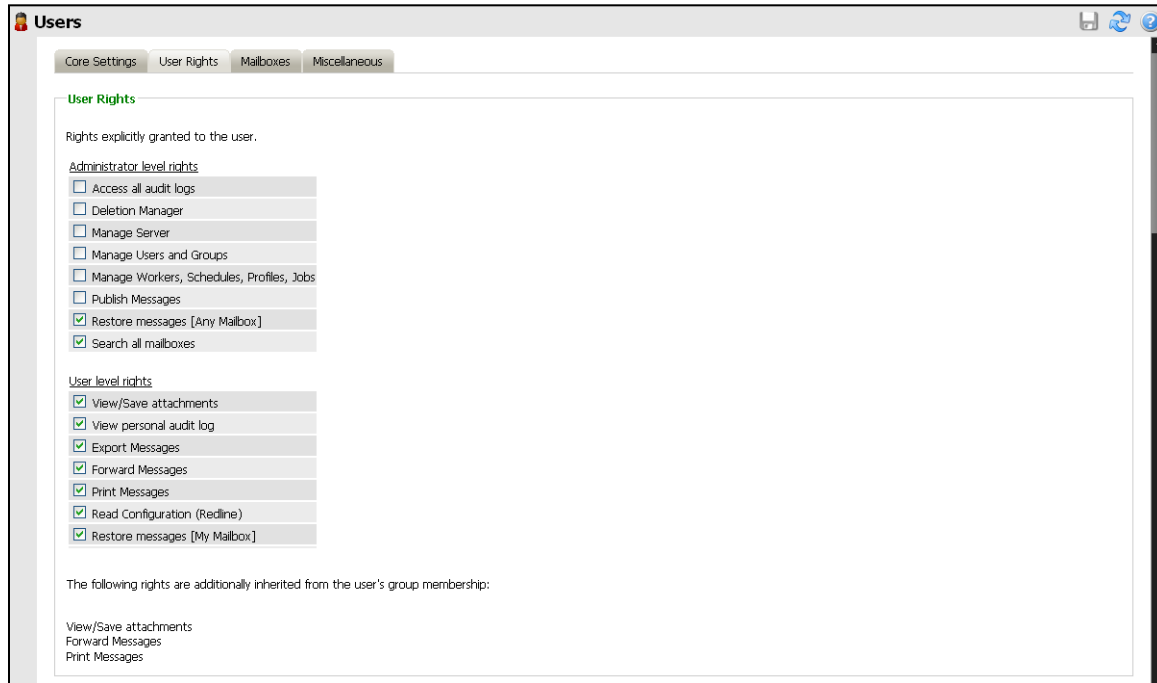
- Authentication Method:** SOAP Authentication (dropdown), with a checkbox for 'Use Exclusively?'.
- Primary UID (links your rights):** (Full rights to all mailboxes) (text field).
- Group Membership:** default (dropdown).
- This account does not expire:** checkbox.
- Change Internal Password:** New (text field), Confirm (text field), and a checkbox for 'Prevent user from changing?'.
- Language:** English (dropdown).
- Disable Account?:** checkbox, with a date field set to 2008 Mar 11.

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.
 - 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
 - Shows the UUID of their SOAP mailbox. All GroupWise users normally have one.
 - The first time GW users log in via SOAP, this field is populated.
 - *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.
- 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
- Search All Mailboxes
 - Grants access to all mailboxes in the system
 - Allows the user to search and browse them.

Administrator level rights	
<input checked="" type="checkbox"/>	Access all audit logs
<input checked="" type="checkbox"/>	Deletion Manager
<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"/>	Search all mailboxes
User level rights	
<input checked="" type="checkbox"/>	View/Save attachments
<input type="checkbox"/>	View personal audit log
<input checked="" type="checkbox"/>	Forward Messages
<input checked="" type="checkbox"/>	Print Messages
<input type="checkbox"/>	Read Configuration (Redline)

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.

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 (You can also view the audit log in the PROPERTIES tab of each message.)
- Forward Messages
- 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 a web application window titled 'Users'. At the top, a status message says 'User colin loaded.' Below it is a text input field for 'User ID' containing the value 'colin'. There are four tabs: 'Core Settings', 'User Rights', 'Mailboxes' (which is selected), and 'Miscellaneous'. The 'Mailboxes' tab contains a section titled 'Additional Mailboxes' with explanatory text. Below this is a 'Currently Selected' list containing two entries: 'Willem Bagchus' and 'Colin James', each with a red 'X' icon to its left. A 'Criteria' section follows, with dropdowns for 'First Name' and 'begins with', a text input field, and a 'Search' button. There are also two checked checkboxes: 'Show only users' and 'Show only recently cached items', and an 'Add Selected Items' button. At the bottom is an 'Address Book' section with a large text area for entering search criteria.

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 used a lot in various other areas, so here is an in-depth discussion.

Currently Selected

✖ Willem Bagchus

✖ Colin James

Criteria

First Name begins with

Show only users ☒ Show only recently cached items ☒

Address Book

	First	Last	PostOffice	Mailbox	Dept.	Type	Domain	UID
<input type="checkbox"/>	Willem	Bagchus	gw2po	Willem		U	gw2dom	109A2120 1112 0000 B003 2C91B3A1EB12
<input type="checkbox"/>		Greg	gw2po	Lynn		U	gw2dom	DA34E130 11C8 0000 9AD4 0D003A008A00
<input type="checkbox"/>	Colin	James	gw2po	Colin		U	gw2dom	4CC19AC0 1112 0000 B003 2C91B3A1EB12
<input type="checkbox"/>		Miss	gw2po	Hanigan		U	gw2dom	69896270 11C9 0000 9AD4 0D003A008A00
<input type="checkbox"/>		admin	gw2po	admin		U	gw2dom	DBDF8400 1104 0000 8EA7 7FEA5E8CB6DF
<input type="checkbox"/>		carl	gw2po	Heinz		U	gw2dom	C9381FF0 11C8 0000 9AD4 0D003A008A00

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) Fill out basic criteria to narrow your search results (or no criteria for the first 100)
- 2) Click **Search**.
- 3) 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.
- 4) Choose which of the results you want to add to the selected list
- 5) 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 GroupWise 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](#).

The screenshot shows the 'Users' management interface with the 'Miscellaneous' tab selected. The interface includes a sidebar with a user icon and the title 'Users'. The main content area has four tabs: 'Core Settings', 'User Rights', 'Mailboxes', and 'Miscellaneous'. The 'Miscellaneous' tab is active, displaying several settings sections:

- Comment:** A section with a description 'The default comment appended to forwarded messages can be set here.' and a text input field labeled 'Comment'.
- Date/Time Formats:** A section with a description 'Defines how dates and times are displayed.' and two dropdown menus: 'Date' (set to 'Day-ShortMonth-Year') and 'Time' (set to '24 Hour'). An example '28-Aug-2008 20:02' is shown below.
- Display Number:** A section with a description 'How many items to display per page' and a dropdown menu labeled 'Display Number' (set to '25').
- Message Age Display:** A section with a description 'How many days of messages should be displayed by default?' and a dropdown menu labeled 'Message Age Display' (set to 'Last 3 days').
- Session Timeout:** A section with a description 'Time in minutes, to expire an inactive login session.' and a text input field labeled 'Time (minutes)' (set to '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 [Server Configuration \(Miscellaneous\)](#) 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.

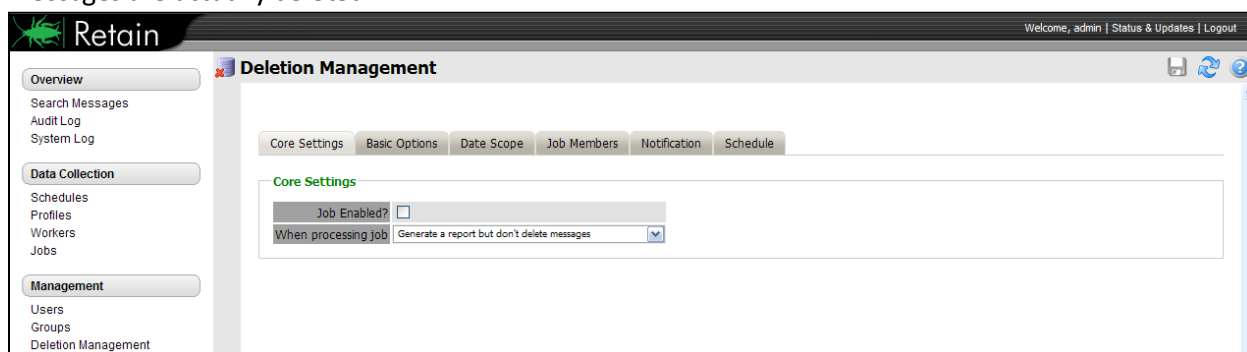
Deletion Management

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 rights. See [User Rights](#).

Core Settings

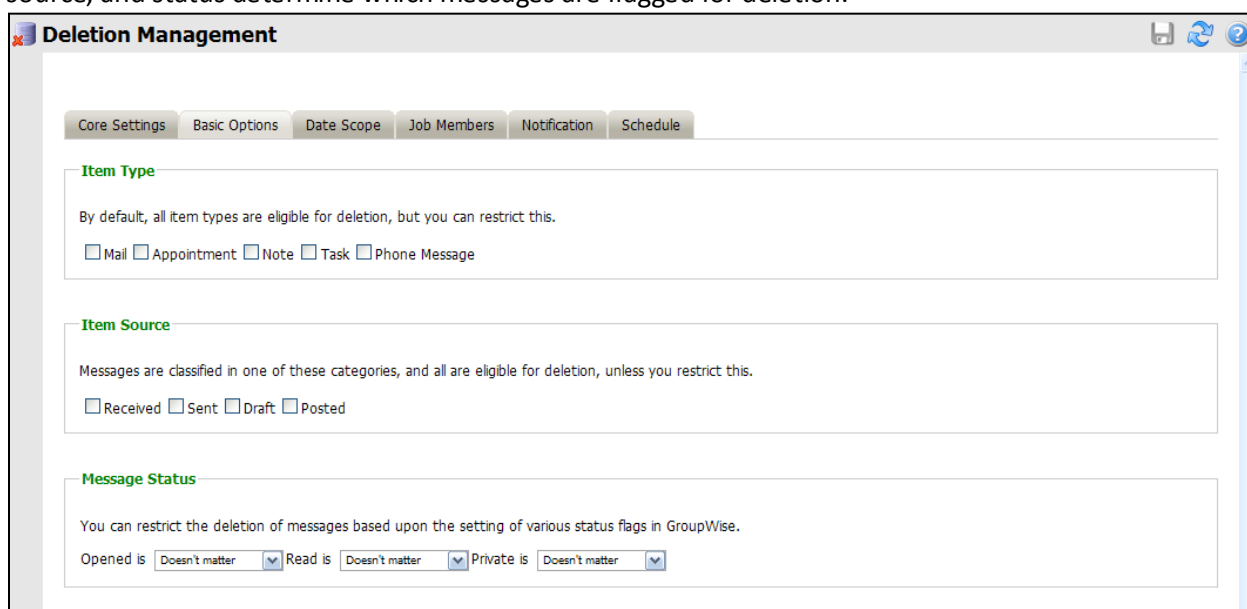
Here you enable and disable deletion jobs, and detail which actions they will take.

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 GroupWise 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.

The screenshot shows the 'Deletion Management' window with the 'Date Scope' tab selected. The interface includes a tab bar with 'Core Settings', 'Basic Options', 'Date Scope', 'Job Members', 'Notification', and 'Schedule'. Below the tabs, a section titled 'Date Scope' contains explanatory text and configuration options. The 'Remove messages using this date filter:' dropdown is set to 'Date Stored in Retain'. The 'Begin Date' and 'End Date' sections each have two options: 'Job Execution Time' (selected) and 'Specific Time'. The 'Job Execution Time' options are set to 'Plus 0 days'. The 'Specific Time' options are set to 'Aug 13 2008' and 'Time: 17:10'. A legend box on the right lists the date types: 'Date Stored in Retain' (highlighted), 'Delivered Date', 'Expiration Date', and 'Created Date'.

Deletion Management

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

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 internal GroupWise dates of the message. The stored date is when the message was stored in Retain. The expiration date is related to the expiration s Job when the message was stored.

Remove messages using this date filter: Date Stored in Retain

Begin Date

☒ Job Execution Time Plus 0 days

☐ Specific Time Date: Aug 13 2008 Time: 17 10

End Date

☒ Job Execution Time Plus 0 days

☐ Specific Time Date: Aug 13 2008 Time: 17 10

Date Stored in Retain
Delivered Date
Expiration Date
Created Date

The dates can be identified by the date filter. The Date filter allows you to specify the GroupWise or Retain message dates. The creation and delivered date are GroupWise 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.

The screenshot shows the 'Job Members' tab in the 'Deletion Management' window. It has a tab bar with 'Core Settings', 'Basic Options', 'Date Scope', 'Job Members', 'Notification', and 'Schedule'. The 'Job Members' section is active. It contains two sections: 'Include these objects:' and 'Exclude these objects:'. The 'Include' section has a dropdown menu with 'po (dom)' selected, an 'Add PO' button, and an 'Add User' button. Below this is a large empty rectangular box. The 'Exclude' section has an 'Add User' button and another empty rectangular box below it.

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.

Job Members

A delete job will only be active for selected users or a selected post office. The Job Members tab allows you to include an entire post office 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. 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.)

The screenshot shows the 'Notification' tab in the 'Deletion Management' window. It has a tab bar with 'Core Settings', 'Basic Options', 'Date Scope', 'Job Members', 'Notification', and 'Schedule'. The 'Notification' section is active. It contains a text box stating: 'You can have Retain e-mail a summary of operations and/or of errors that occurred when the job completes.' Below this are several form fields: 'Mail when Errors Occur?' with a checkbox, 'Mail summary when job complete?' with a checkbox, 'SMTP To Address' (empty), 'SMTP Mail From Address' (Retain@bricebitter.co), 'SMTP Mail Server' (10.1.1.101), 'SMTP Username' (chris), and 'SMTP Password' (masked with dots). There is a 'Test Connection' button at the bottom.

The screenshot shows the 'Schedule' tab in the 'Deletion Management' window. It has a tab bar with 'Core Settings', 'Basic Options', 'Date Scope', 'Job Members', 'Notification', and 'Schedule'. The 'Schedule' section is active. It contains a text box stating: 'A deletion job executes during the maintenance period on a weekly or monthly basis. A deletion job can take a very long time to run, and be quite intensive, if you have many messages to delete.' Below this is a 'Run Job when' section with two dropdown menus: the first is set to 'weekly' and the second is set to 'Sunday'.

Configuration

Server Configuration

The overall operation of Retain is configured in this section.

The screenshot shows the Retain web application's 'Server Configuration' page. The interface has a dark header with the Retain logo and a user status bar. A left sidebar contains navigation links for Overview, Data Collection, Management, and Configuration. The main content area is titled 'Server Configuration' and features a tabbed interface with tabs for Communications, Index, Storage, Accounts, Maintenance, Notification, Logging, and Miscellaneous. The 'Communications' tab is active, showing three sections: 'Database Connectivity', 'Trusted Application Key Generation', and 'GroupWise SOAP Access'. Each section contains descriptive text, instructions, and input fields for configuration.

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 retain created. Be sure to specify:

- The correct JDBC URL. If the database server is not running on the same machine as Retain, 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.

Database Driver:

Database URL:

Database User Account:

Database User Password:

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.

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

Trusted Application Key:

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:

SOAP Port:

Enable SSL? ☐

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

Topics covered on the communications page are covered elsewhere in this manual – look for [initial setup](#) of Retain.

- [Database Connectivity](#)
- [Trusted Application Key](#)
- [GroupWise SOAP Access](#)

Index



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), OpenOffice2, DOC, and PDF. (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.) Add as many as you need.

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.

Communications Index Storage Accounts Maintenance Notification Logging Miscellaneous

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.

<input checked="" type="checkbox"/> Extension	Item	Stream Size	File Size	
<input checked="" type="checkbox"/>	log	204800	1048576	Delete
<input checked="" type="checkbox"/>	eml	204800	1048576	Delete
<input checked="" type="checkbox"/>	docx	2048576	2048576	Delete
<input checked="" type="checkbox"/>	html	204800	1048576	Delete
<input checked="" type="checkbox"/>	cfg	204800	1048576	Delete
<input checked="" type="checkbox"/>	xml	204800	1048576	Delete
<input checked="" type="checkbox"/>	rtf	204800	1048576	Delete
<input checked="" type="checkbox"/>	txt	204800	1048576	Delete
<input checked="" type="checkbox"/>	odt	2048576	2048576	Delete
<input checked="" type="checkbox"/>	odp	2048576	2048576	Delete
<input checked="" type="checkbox"/>	doc	2048576	2048576	Delete
<input checked="" type="checkbox"/>	ods	2048576	2048576	Delete
<input checked="" type="checkbox"/>	odg	2048576	2048576	Delete

[Force Indexing Now](#)

Storage

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

Storage

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

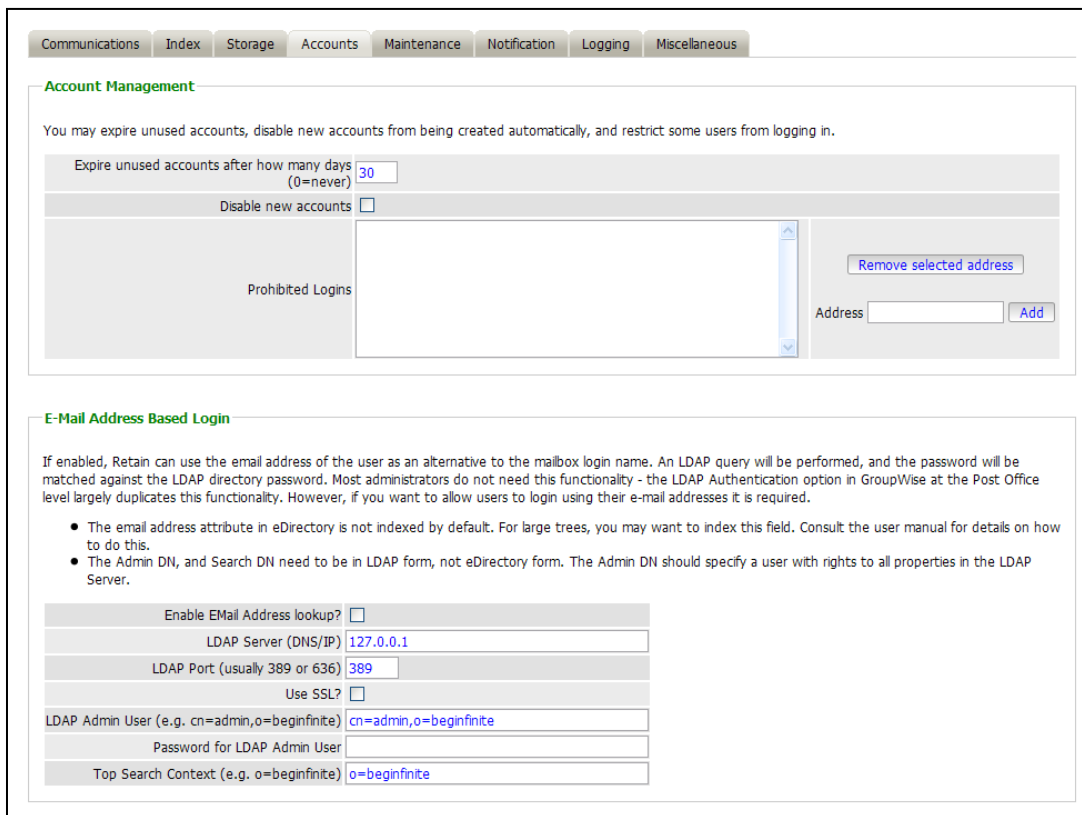
Encryption

Compression (always on)

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



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](#)

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)

LDAP Port (usually 389 or 636)

Use SSL? ☐

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

Password for LDAP Admin User

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

Open System vs. Closed System

Normally, Retain lets all GroupWise 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.

LDAP and E-Mail Address Based Login

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?	<input type="checkbox"/>
LDAP Server (DNS/IP)	127.0.0.1
LDAP Port (usually 389 or 636)	389
Use SSL?	<input type="checkbox"/>
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

To use LDAP authentication for [users to login](#) to Retain, this section must first be filled-out correctly. Retain needs to know where the LDAP Server is, which port it is listening on, and the login information for the administrator so that e-mail addresses can be checked.

E-mail address based logins can take longer in larger systems, if the e-mail addresses are not cached by the LDAP server. By default, e-mail addresses are not cached by the LDAP server, so you must set them to be cached. This setting is configured in ConsoleOne.

To enable email address caching:

In *ConsoleOne*, right-click on the *server* object and select *Properties*.
Scroll right until you see the *'Indexes'* tab.
Select *'Add'*

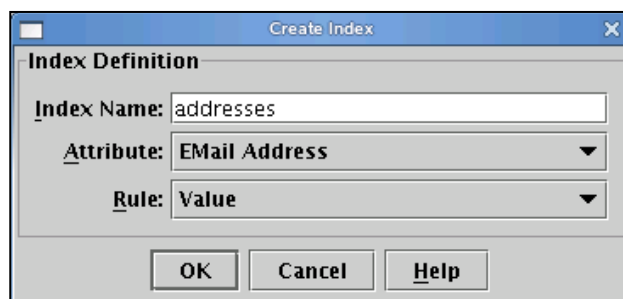
In the Create Index window, name the index what you desire. (Since this will be caching e-mail addresses, an appropriate name may help in the future.)

From the Attribute list, scroll down and select *'Email Address'*.

Set the Rule to *'Value'*.

Click *'OK'*.

Email Addresses will now be cached.



Maintenance

Maintaining a Retain system involves many tasks:

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

These items are all controlled here.

Maintenance/Caching

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.

Retain scans the GroupWise address book on a regular interval, and permanently stores user, Post Office, Domain, and Distribution List information. This is stored in a local database for Retain to be able to access at all times. It is critical this address book be kept up to date.

Enable Backup of Embedded Database	<input checked="" type="checkbox"/>	Sunday	weekly
Enable Backup of Indices	<input checked="" type="checkbox"/>	Sunday	weekly
Enable Index Optimization	<input checked="" type="checkbox"/>	(every day)	weekly
Enable Compression of Embedded Database	<input checked="" type="checkbox"/>	(every day)	
Enable Address Book Caching	<input checked="" type="checkbox"/>	(every day)	
Run maintain procedure at	01	00	
Last Backup of DB	(never)		
Last Backup of Indices	(never)		
Last Index Optimization	(never)		
Last Compression	(never)		
Last Address Book Cache	05-Mar-2008 11:57:29		

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.

The system address book is cached from the GroupWise address book. It is read and updated daily by default. Here you configure whether or not to automatically do this and when.

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)

Post Office Redirection Tables

Retain has overrides for the addresses of known post offices. While Novell's SOAP protocol furnishes addresses to the post offices, sometimes they are not correct. For example, the IP addresses may not be reachable by the Retain Worker across a NAT or WAN link. So you may manually enter an address or DNS name that the Retain Worker will use to access the selected Post Office Agent.

A classic example of this is the case of GWAVA Reload. If you wish to draw the data from a Reload server, you will use the same Post Office Agent name but the SOAP protocol will furnish the address of the LIVE POA. In this case, you would put the address/DNS name of the Reload POA so the Retain Worker will draw the data from that server instead of the live one.

The login column allows you to specify which server to use for AUTHENTICATION purposes – for instance, when a user logs in to Retain.

The Worker column is used to specify which server the WORKER should use to draw data from.

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.

[Refresh Address Book](#)

Post Offices	Login (Force by selecting checkbox)	Worker (Force by selecting checkbox)
gw2po (gw2dom)	192.168.100.17 : 7191 <input checked="" type="checkbox"/>	192.168.100.17 : 7191 <input checked="" type="checkbox"/>

Click “Refresh Address Book” to update Retain on your GroupWise system and in particular, to make it aware of any changes to your GroupWise system or users.

Ordinarily, Retain refreshes this every night but you can force a refresh here.

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

[Remove Selected Domain](#)

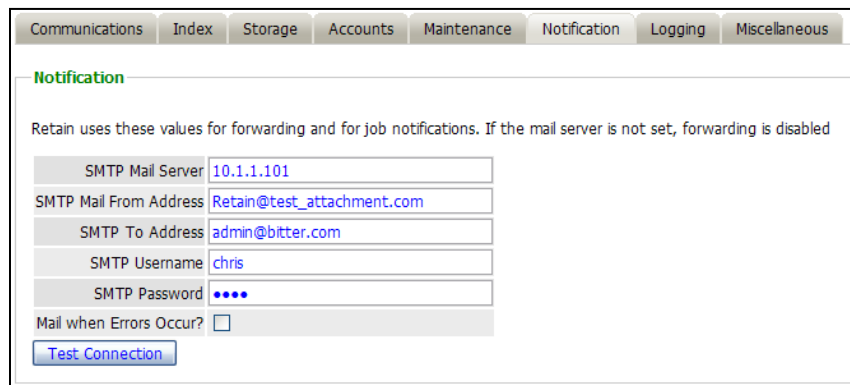
Domain [Add](#)

GroupWise allows you to add non-GroupWise (external) domains to the system. Sometimes, you might do this to add external domains to the GroupWise address book. However, GroupWise 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.

Notification

On this screen, you can configure the SMTP options so that Retain may send e-mail for notification or forwarding from the search screen.



The screenshot shows the 'Notification' configuration screen in the GWAVA Retain application. At the top, there is a navigation bar with tabs: Communications, Index, Storage, Accounts, Maintenance, Notification (selected), Logging, and Miscellaneous. Below the tabs, the title 'Notification' is displayed in green. A note states: 'Retain uses these values for forwarding and for job notifications. If the mail server is not set, forwarding is disabled'. The configuration fields are as follows:

SMTP Mail Server	10.1.1.101
SMTP Mail From Address	Retain@test_attachment.com
SMTP To Address	admin@bitter.com
SMTP Username	chris
SMTP Password	••••
Mail when Errors Occur?	<input type="checkbox"/>
Test Connection	

Retain logs in to the SMTP server (usually a GWIA) to authenticate before sending a message. You may also test the setup to be sure it works.

This is another step you complete during [initial setup](#) and you may change it here as needed.

Logging

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

Communications	Index	Storage	Accounts	Maintenance	Notification	Logging	Miscellaneous
Logging							
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.

Auditing
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) <input type="text" value="0"/>
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"/> Server Config Change
<input type="checkbox"/> Audited

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	Notification	Logging	Miscellaneous
Miscellaneous							
<p>You may change the default Worker password, which is assigned to all Workers when they are created (but can be individually overridden).</p> <p>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.</p> <p>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.</p>							
Default Worker Password		<input type="password" value="....."/>					
Export XML		<input type="checkbox"/>					
Enable support for GroupWise Proxy		<input type="checkbox"/>					
Cache Proxy verifications (days)		<input type="text" value="7"/>					

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.

GroupWise Proxy support is also enabled here. Just check the box and specify the refresh interval. (Cache Proxy verifications) Since it is a very performance intensive operation, Retain will reload and recheck proxy settings according to the interval you specify. The default is every 7 days.

Export XML

This XML Export function is included in Retain in case you have an XML compliance mandate. You enable it here.

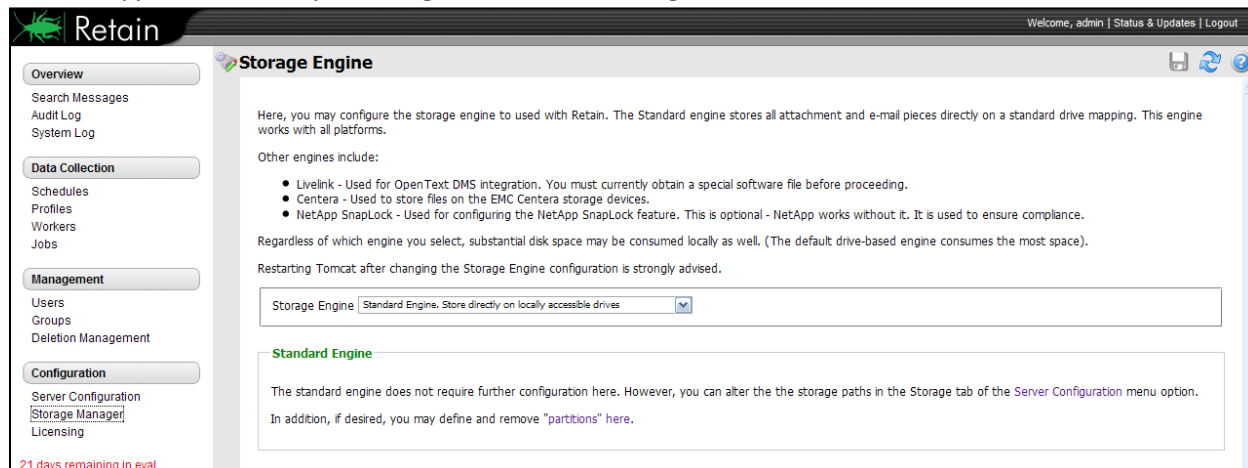
When selected, each attachment will have an XML export file of its parent message. In other words, an XML representation of the metadata is created and linked to the blobs as the messages are stored. There's nothing more done.

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.

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.

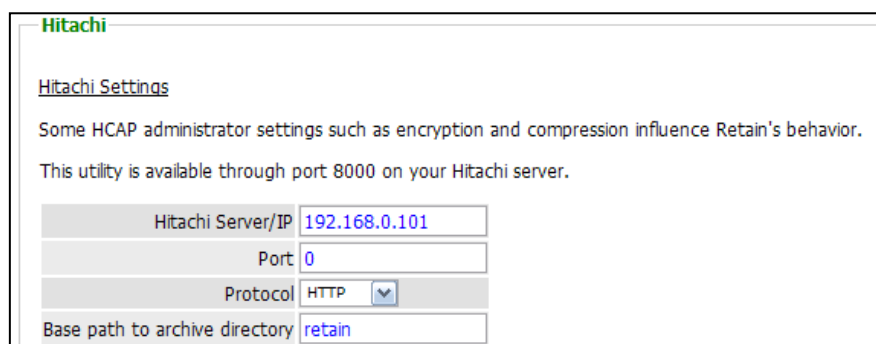
If the Standard Engine is selected, then the Standard Engine window is displayed. This details the setting location for the default storage location, under Server Configuration, Storage tab.

The other four options are:



Hitachi Content Archive Platform

The Hitachi Content Archive Platform must allow HTTP or HTTPS access from the Stub Server IP address. Check these settings on the HCAP Administration console. HTTP is recommended for performance.



Retain needs the following information;

- Hitachi Server/IP - DNS is preferred because IP addresses can quickly change, both are acceptable.
- Port - the port number where HCAP is listening. (Port 80 or 0)
- Protocol - Specify HTTP or HTTPS, which ever protocol your system is using.
- Base path to the archive directory – Specify where Retain will place archived data. **This path forms a URL on a web server and is NOT a mapped drive.** Use the following syntax:
 <directory>/ sub directory (i.e. Retain/data)
 If the Base Path does not yet exist, Retain will create it with the first running job.

Livelihood / Open Text DMS System

Because Livelihood requires each user to exist in the Livelihood system, Retain now automatically creates users in Livelihood as their mail is archived. In addition, Livelihood requires the following information to allow you to use the system.

Server Connectivity Settings

- *Livelihood Server/IP* - This should be set to the DNS or IP address of the Livelihood system. DNS is preferred because IP addresses can quickly change.
- *Port* - The default port for the LAPI, (Livelihood API), is port 2099, unless you are using LAPI over HTTP, which requires port 80.
- *User / Password* - The user required by Retain needs the rights to 'impersonate', or login to all the users which are going to be archived by Retain. If you are going to utilize user auto creation, then you will need the right to create users. This can all be done by a system administrator account, and does not need to be the super administrator.
- *URL Path* - The URL path is only required if you use LAPI over HTTP. The path shown is the default for Windows. The UNIX default path is /livelihood/livelihood. If you are not using LAPI over HTTP, then this information is not required or used.
- *Use LAPI over HTTP* - LAPI over HTTP is suggested and may be required by Livelihood support, but it is slower than a raw connection to the LAPI on port 2099. If you set this to true, you must also change the connection port to 80 and verify the correct URL path.

The screenshot shows the 'Livelihood' configuration window. It is divided into three sections: 'Server Connectivity Settings', 'Localization settings', and 'Optional Settings'.
 - **Server Connectivity Settings:** Includes fields for 'Livelihood Server/IP', 'Port' (set to 2099), 'User', 'Password', and 'URL Path' (set to /livelihood/livelihood.exe). There is a checkbox for 'Use LAPI over HTTP' which is currently set to 'False'.
 - **Localization settings:** Includes a checkbox for 'Enable folder translation?' (set to 'False'), a checkbox for 'Prepend Headers to Body Text?' (set to 'True'), and a 'Language' dropdown menu set to 'English'.
 - **Optional Settings:** Includes checkboxes for 'Set profile name as category?' (set to 'False'), 'Set profile name as comment?' (set to 'False'), and a 'Debug Proxy' dropdown menu set to 'False'.

Localization Settings

- *Enable folder translation?* - Retain will create folders in Livelihood that match the folder names that Retain receives from the Post Office. Folder translation will change the folder system names to the equivalent name in the language selected below. If this is set to false, then the bare name received will be the name set in the Livelihood system.
- *Prepend headers to Body Text?* - Retain can add five lines to the top of each message that contain the: from, to, cc, bcc, subject, and date. This option should be set to true, as it is a good idea to keep this information with each file, but is allowed to be switched-off here for legal compliance. If you are legally prohibited from altering the mail files in any way, then you should turn this option off. Otherwise leave it on.
- *Language* - You can set the local language option here. This option is useless if the above folder translation option is set to false. The supported languages are English, German, French, and Spanish.

Optional Settings

- *Set Profile name as category? And Set profile name as comment?* - Livelihood has the ability to manage files according to a Livelihood profile which can use a comment or category placed on the files stored. To use this, enable the option desire here, either comment or category, and the category specified by Retain will be the name of the active [Profile](#).
- *Debug Proxy* - The debug proxy is a diagnostic tool only. This option will log all data transferred to Livelihood and will create HUGE data logs. DO NOT enable this option unless specifically instructed to do so by Support.

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	<input type="text" value="80"/>
Admin User	<input type="text"/>
Admin Password	<input type="text"/>
Base path to archive directory (example: /vol/myvolname/archive)	<input type="text"/>

Partitions

In addition, if desired, you may define and remove "partitions" here.

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

Begin Date	Name	Path	Delete
04-Mar-2008 12:53:24	<input type="text" value="new partition"/>	<input type="text" value="/var/opt/beginfinite/retain/new"/>	<input type="button" value="Delete"/>

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 Standard Engine window. (Shown below) 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)

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

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

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 consulting support first.**

If a partition is deleted Retain will

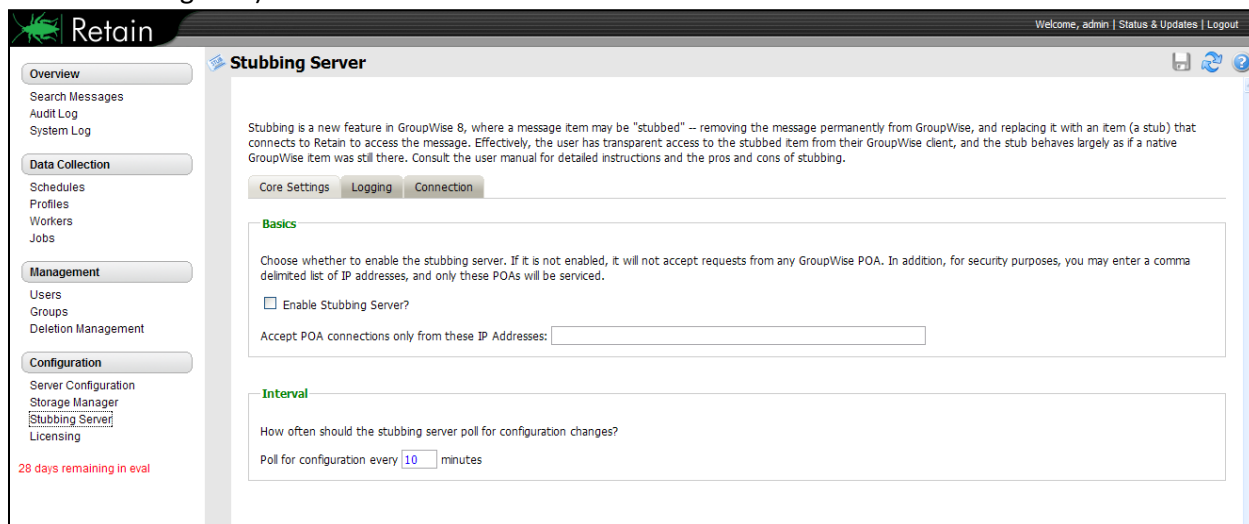
look for the archived mail that was stored in the deleted partition to be stored in the default storage location. (This is defined in Server Configuration, under the [Storage tab](#).) Removing a designated partition in the management console does not delete the storage directories or data on the hard drive or SAN, but it does disconnect Retain from the storage location.

All partition changes are logged in the partition.log file, located in the backup directory specified in Server Configuration. (By default this is located in <your storage location>/backup)

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

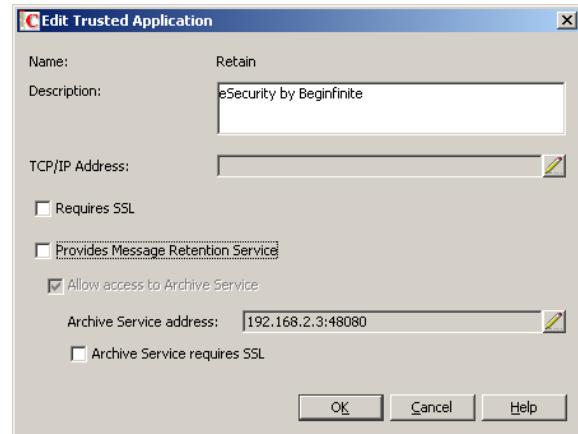
For the Retain Server to function, GroupWise must be told where the Stub Server is located. In ConsoleOne;

- 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"



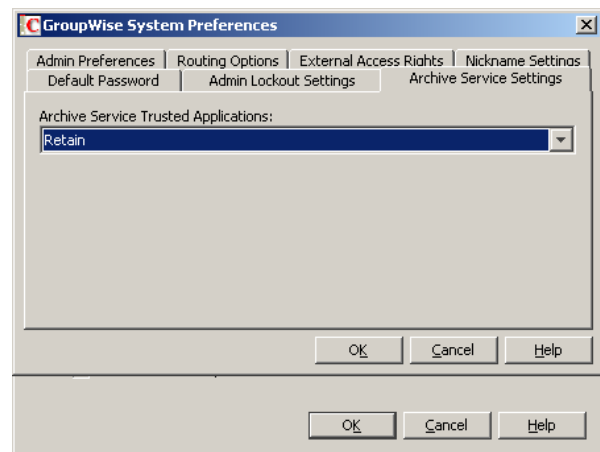
- Put the address or DNS name of the Stub Server and the correct port as shown. (48080)
- Click OK and exit the rest of the windows.

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

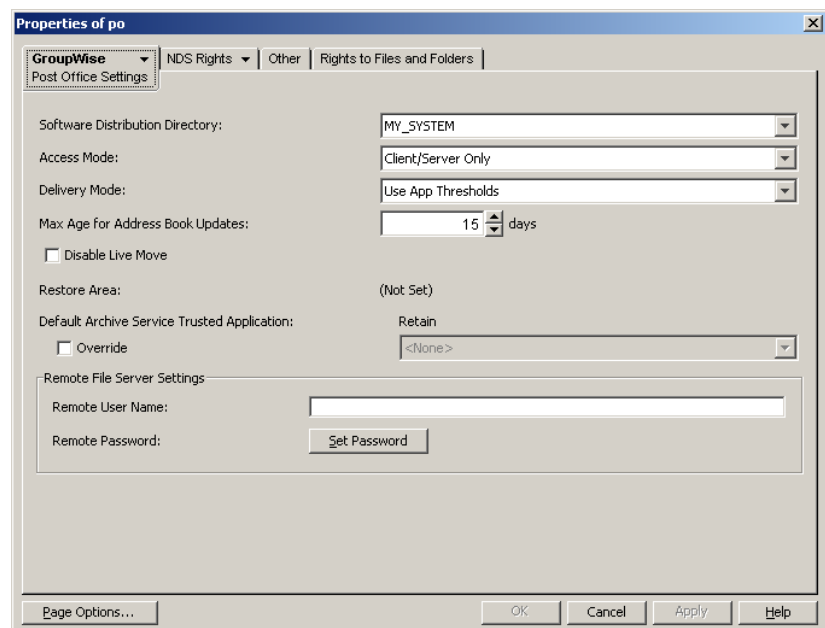


We also need to change the settings in the preferences menu:

- 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 agent. 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' configuration window. It has four tabs: 'Core Settings', 'Logging', 'Connection', and 'Bootstrap'. The 'Core Settings' tab is active, showing two sections: 'Basics' and 'Interval'.

Basics

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.

☒ Enable Stubbing Server?

Accept POA connections only from these IP Addresses:

Interval

How often should the stubbing server poll for configuration changes?

Poll for configuration every 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' configuration window with the 'Logging' tab active. The 'Logging Options' section is visible.

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)	<input type="text" value="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' configuration window with the 'Connection' tab active. The 'Retain Server Connection' section is visible.

Retain Server Connection

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.

Server Protocol	http
Retain Server Host Name	192.168.0.101
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.



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>


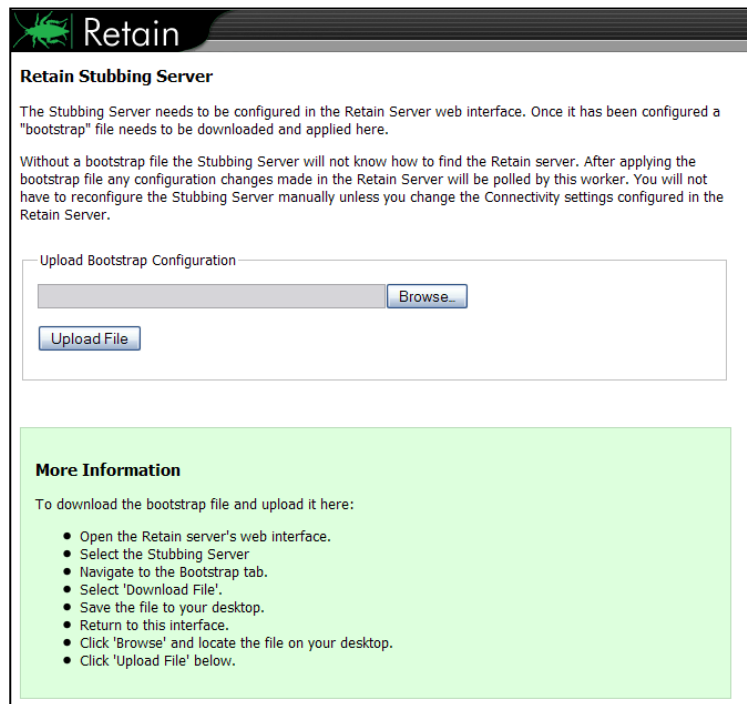
You should see a window like the one pictured here.

- 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

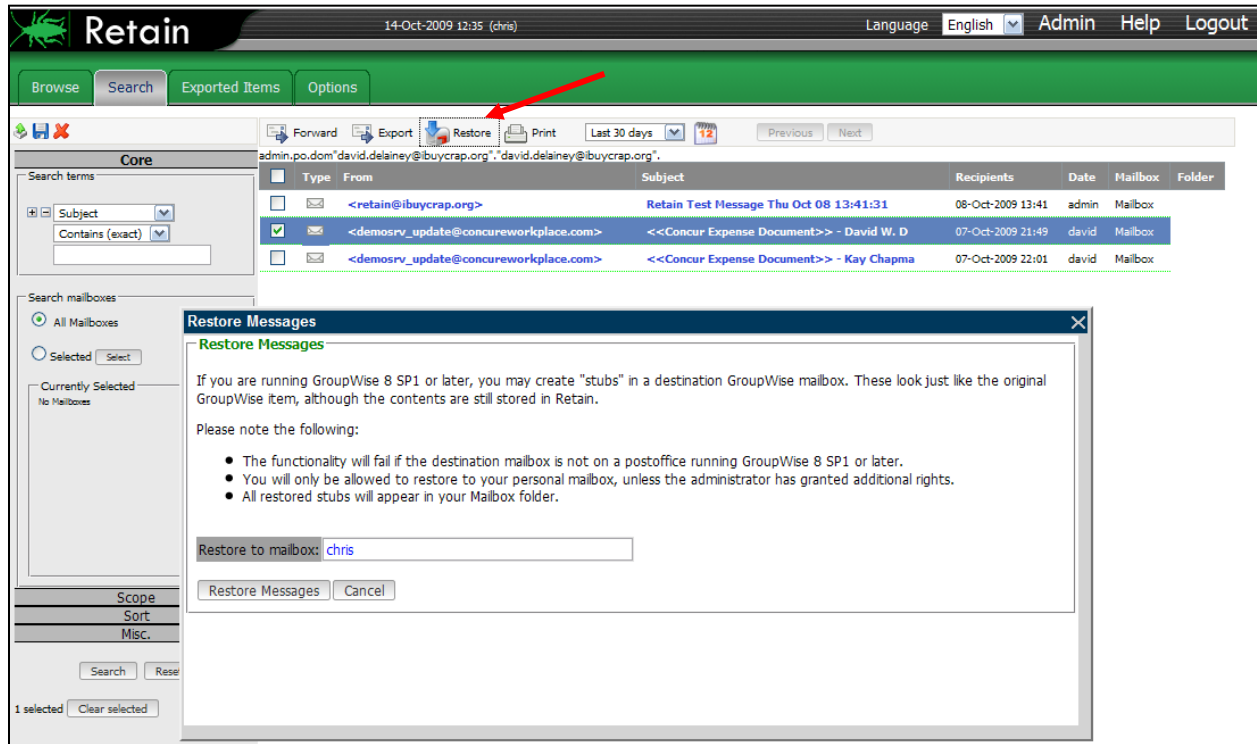
Scheduler

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](#).'

Restore Stub functionality

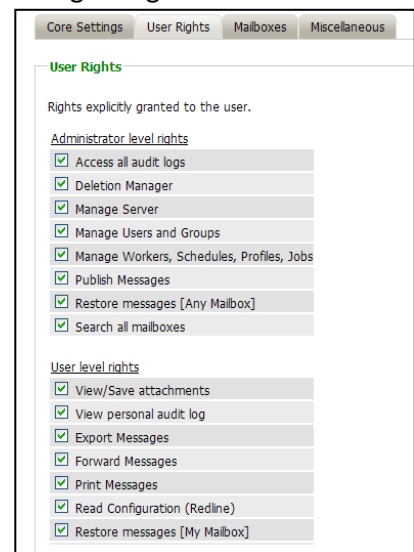
Retain can now restore a message, in stub form, to a mailbox where it has been deleted. Stubbed messages which have been deleted from the GroupWise system from the GroupWise client can be restored to the mailbox from within Retain. Users may, if they have been granted user rights to do so, log into their Retain archive, locate and select the desired message, or messages, and click on the 'Restore' button. 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.



After the 'Restore' option has been selected, the target mailbox must be specified, after which Retain connects to the GroupWise system and places a stub of the selected message in the target mailbox. The Users Right that is required to use this new feature is the "Restore messages" right. 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.

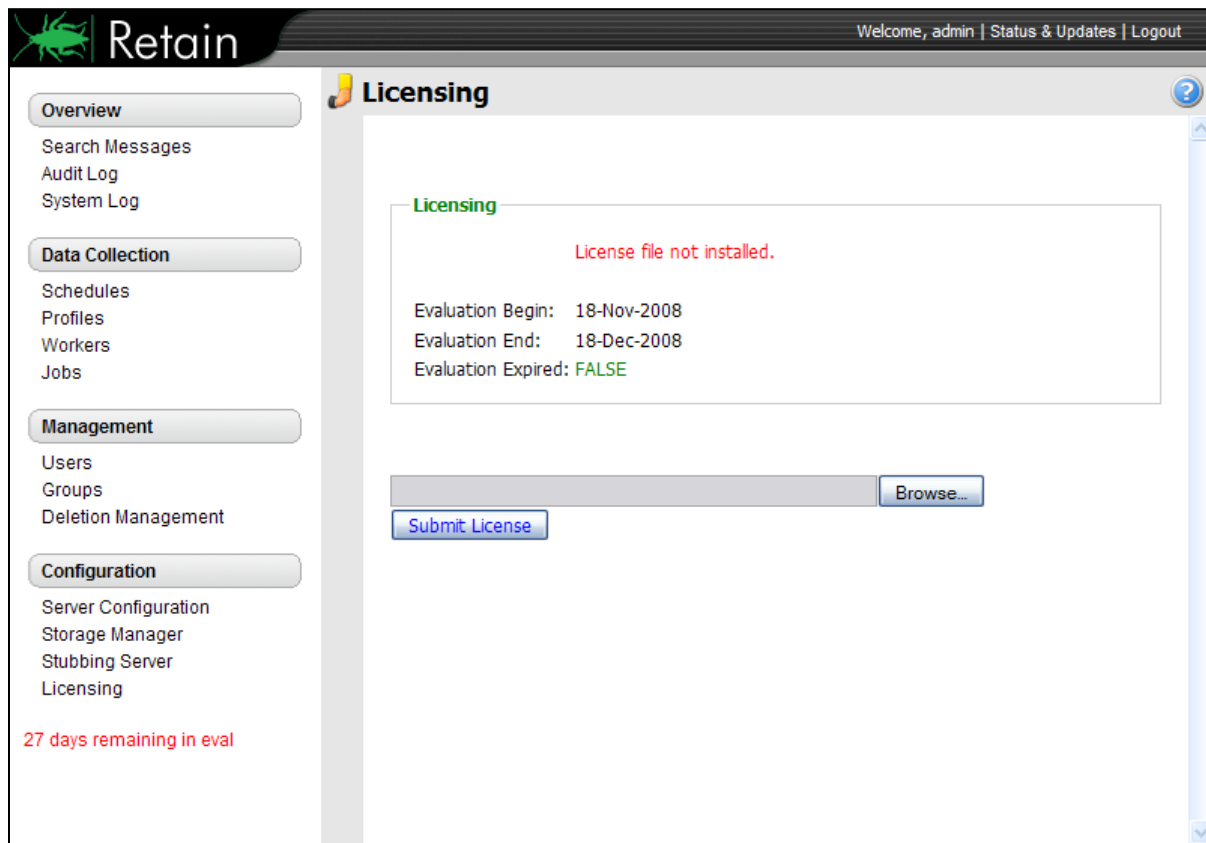
The Restore Stub feature requires:

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



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](#).



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".

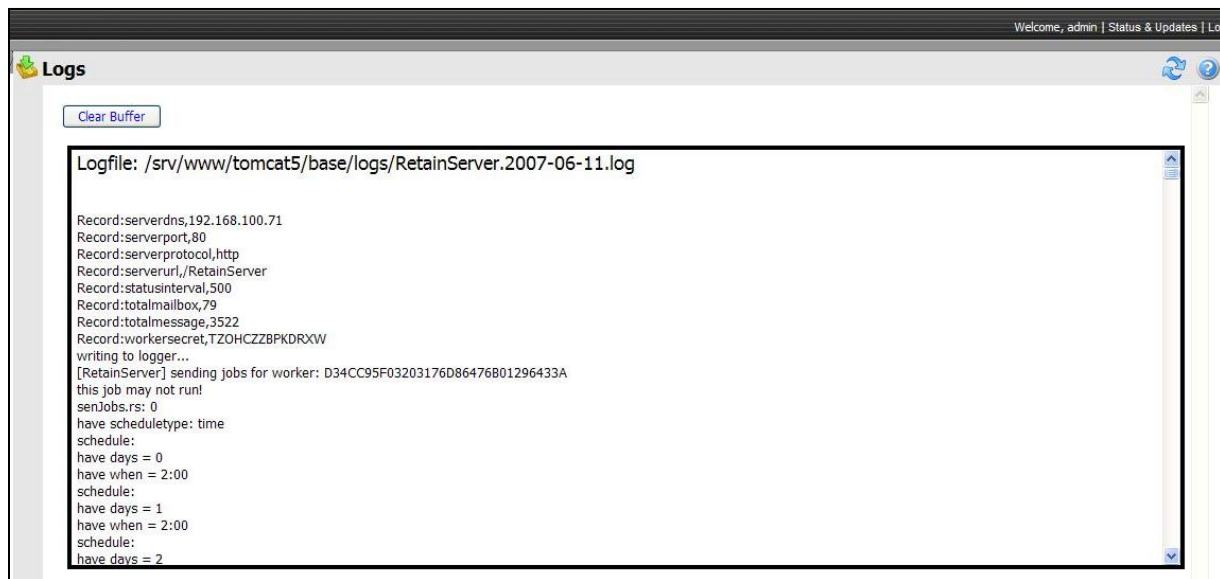
Logout and re-login to the server to see the license in effect.

Overview

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.



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.

Date	Info	User	Mailbox	IP Address	Action
08-Jun-2007 03:08	Agent Accounting Data File	linuxworker1	taylor	192.168.100.71	Archived
08-Jun-2007 03:08	Agent Accounting Data File	linuxworker1	taylor	192.168.100.71	Archived
08-Jun-2007 02:03	Retain Error	linuxworker1	willem	192.168.100.71	Archived
08-Jun-2007 02:03	Retain Error	linuxworker1	willem	192.168.100.71	Archived
08-Jun-2007 02:03	Retain Error	linuxworker1	willem	192.168.100.71	Archived
08-Jun-2007 02:03	Retain Error	linuxworker1	willem	192.168.100.71	Archived
08-Jun-2007 02:03	Retain Status Information	linuxworker1	willem	192.168.100.71	Archived
07-Jun-2007 14:09	Retain Test Message Wed Apr 25 17:06:53 EDT 2007.txt	admin	Willem	192.168.100.126	Viewed
07-Jun-2007 14:05	Retain Test Message Wed Apr 25 17:06:53 EDT 2007.txt	admin	Willem	192.168.100.126	Viewed
06-Jun-2007 02:48	Retain Test Message Tue Jun 05 09:30:33 EDT 2007	linuxworker1	willem	192.168.100.71	Archived
06-Jun-2007 02:48	Retain Status Information	linuxworker1	willem	192.168.100.71	Archived
06-Jun-2007 02:45	Agent Accounting Data File	linuxworker1	taylor	192.168.100.71	Archived
05-Jun-2007 11:44	Retain Test Message Wed Apr 25 17:06:53 EDT 2007.txt	guest	taylor	192.168.100.126	Viewed
05-Jun-2007 11:35	Retain Test Message Wed Apr 25 17:06:53 EDT 2007.txt	guest	taylor	192.168.100.126	Viewed
05-Jun-2007 02:49	Retain Status Information	linuxworker1	willem	192.168.100.71	Archived
05-Jun-2007 02:47	Agent Accounting Data File	linuxworker1	taylor	192.168.100.71	Archived

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.

Searching the Retain Archives

Now that you have data in your Retain archives, you will want to retrieve this information. This is the heart and soul of the information retrieval system in Retain. You search for what you want, the results appear and you can view, print 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.

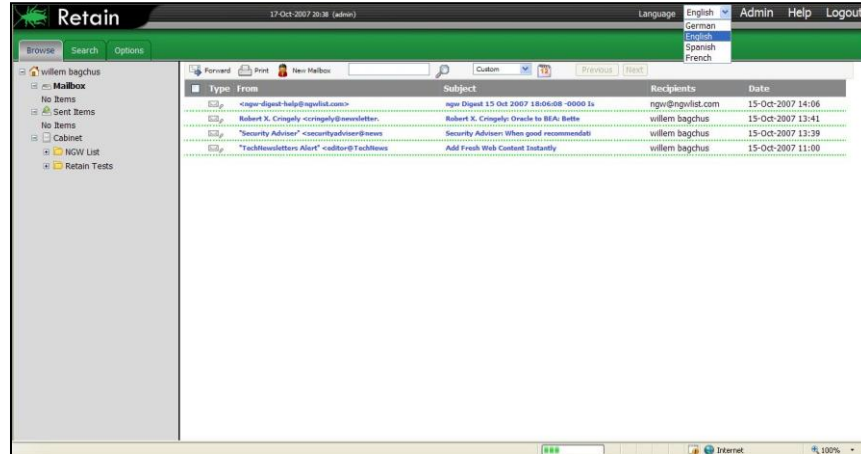
There are three tabs in the Search Interface:

- Browse- for casual browsing and filtering. You may only browse one mailbox at a time.
- Search- for in-depth search, possibly across many mailboxes.
- 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 GroupWise. The screen is designed to resemble GroupWise WebAccess so that users may become quickly familiar with it.

If you have GroupWise Proxy rights, and the proxy function is enabled or if you have the rights to access more than one mailbox, 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.

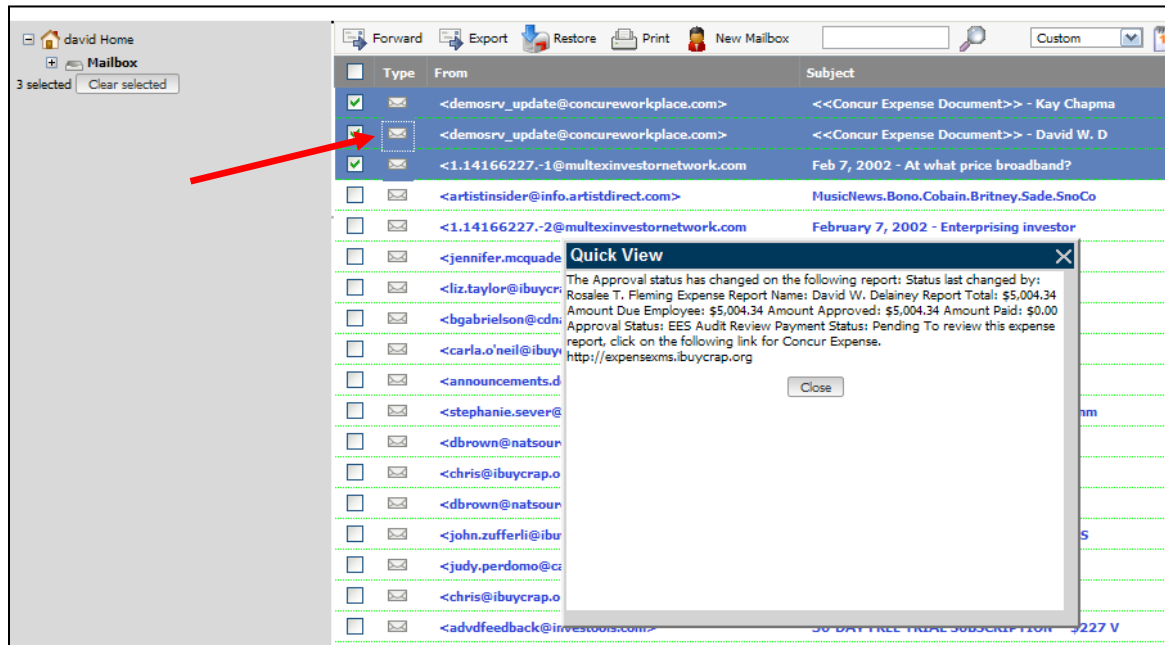


Cumulative Selections

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.



Browsing Tools (Widgets)



Retain will display various options for the user. The specific options displayed depend upon the user's rights. These include:

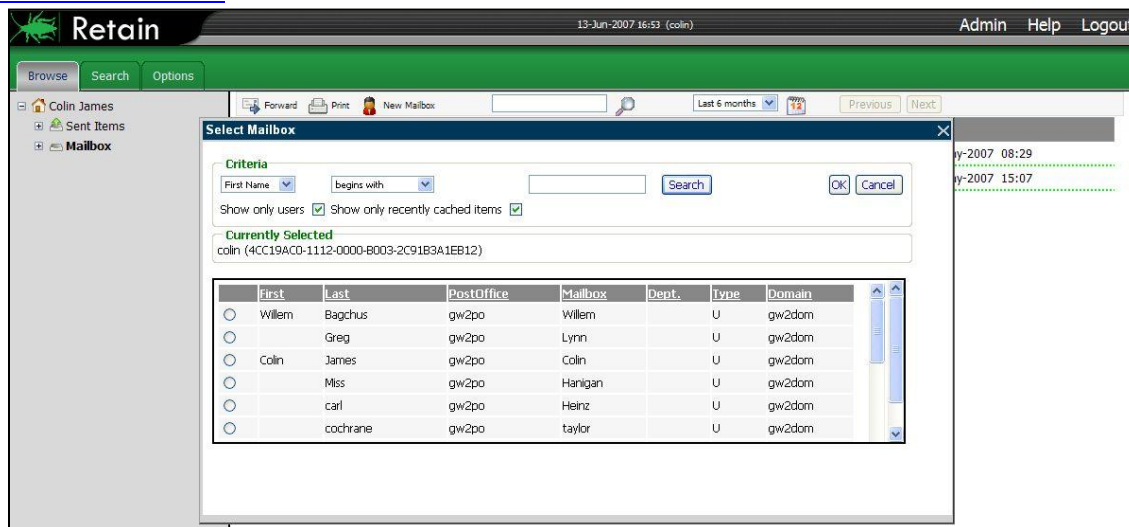
- **FORWARD:** forward selected items by e-mail.
- **PRINT:** prints a list of selected items.
- **NEW MAILBOX:** select a new mailbox to view.
- **QUICK SEARCH:** search the current folder for text in “from”, “to”, “subject” or body.
- **DATE SELECTOR:** select a custom date range. Retain will show 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.

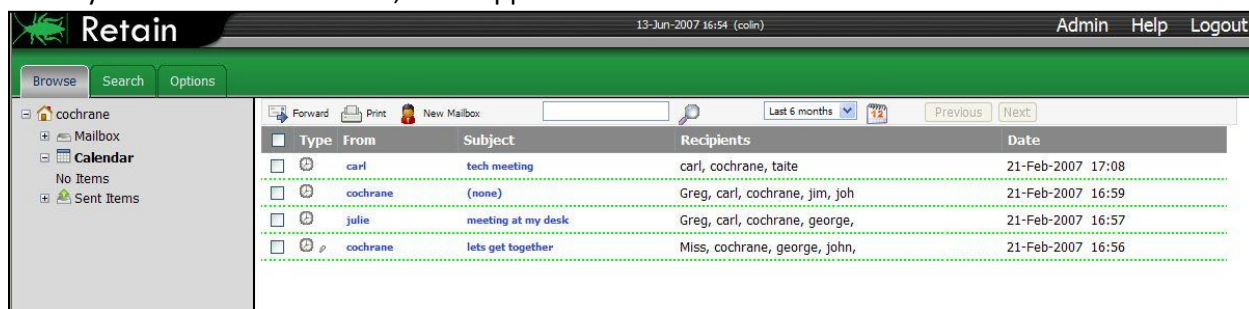
New Mailbox

If you click “New 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](#).



In Retain, you can only browse one mailbox at a time. To search multiple mailboxes simultaneously, use the search function.

When you view another mailbox, it will appear like this:



Notice how calendar items appear.

Toolbar Options:

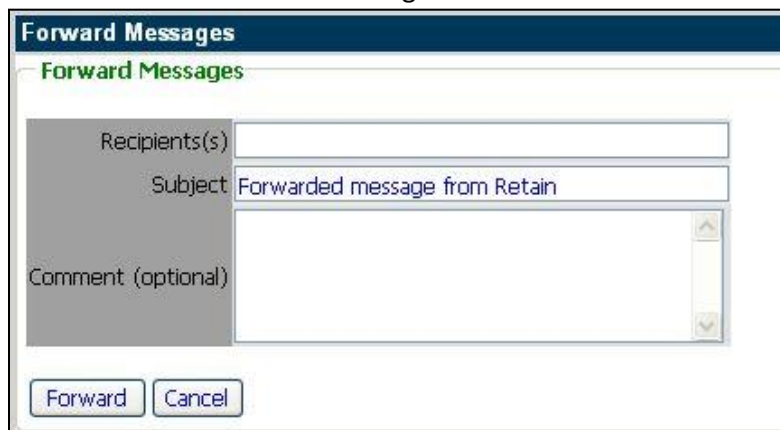
Specifying a Date Range

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



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:

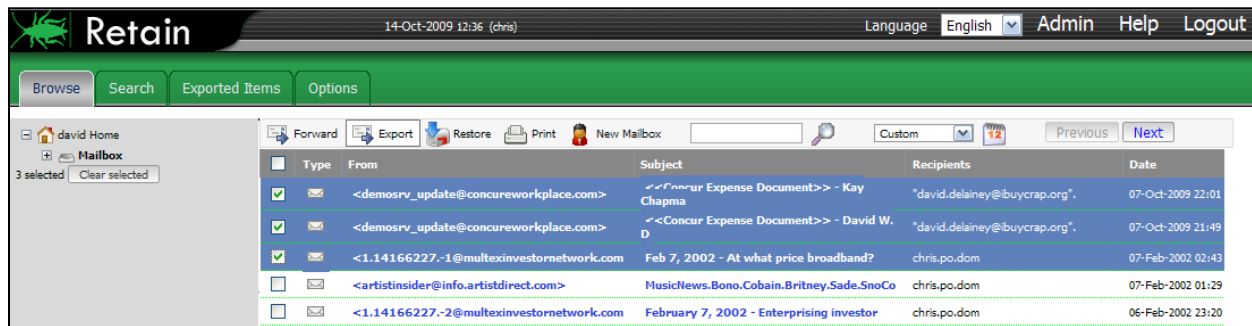


Enter the recipients (where you want the message(s) to be forwarded), separated by commas. Enter a comment if desired and click "**Forward**". Retain will send the message(s) via SMTP.

PDF Export

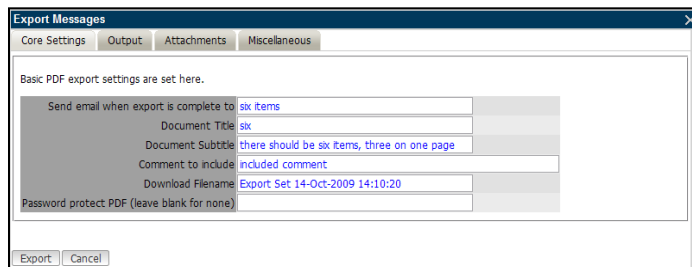
PDF Export is a new function which allows users, (with appropriate rights), to export single or multiple items into a PDF file. This PDF can be named and defined by the user exporting the document. Items will appear as navigable files in the PDF. The default view provides a list on top of the PDF viewer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

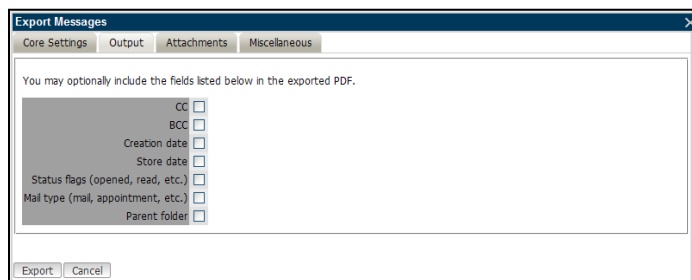


Fill out the various options as desired.

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)	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.

The screenshot shows the Adobe Reader interface. The top pane displays a table of 100 items, which are email messages. The bottom pane shows the selected item's title page, which includes the Retain logo, a title, subtitle, sender, date, and a comments section.

Date	Subject	From	To	Size
9/18/2001 4:20:52 PM	New Contact Info.	"chris@ibuycrap.org".GWIA.dom	"bgabrielson@cdnam.com".GWIA.dom	2.658 KB
9/18/2001 4:20:52 PM	New Contact Info.	"chris@ibuycrap.org".GWIA.dom	"bgabrielson@cdnam.com".GWIA.dom	2.658 KB
9/18/2001 4:20:52 PM	New Contact Info.	"chris@ibuycrap.org".GWIA.dom	"bgabrielson@cdnam.com".GWIA.dom	2.659 KB
9/18/2001 4:20:52 PM	New Contact Info.	"chris@ibuycrap.org".GWIA.dom	"bgabrielson@cdnam.com".GWIA.dom	2.659 KB
12/5/2001 7:36:59 AM	FW: New Power	delaney.GWIA.dom	"rick.buy@ibuycrap.org".GWIA.dom	2.889 KB

Cover sheet

 **Retain**

Title

Subtitle


admin

Thu Dec 03 11:45:53 MST 2009

Comments:
included comment

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.

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.

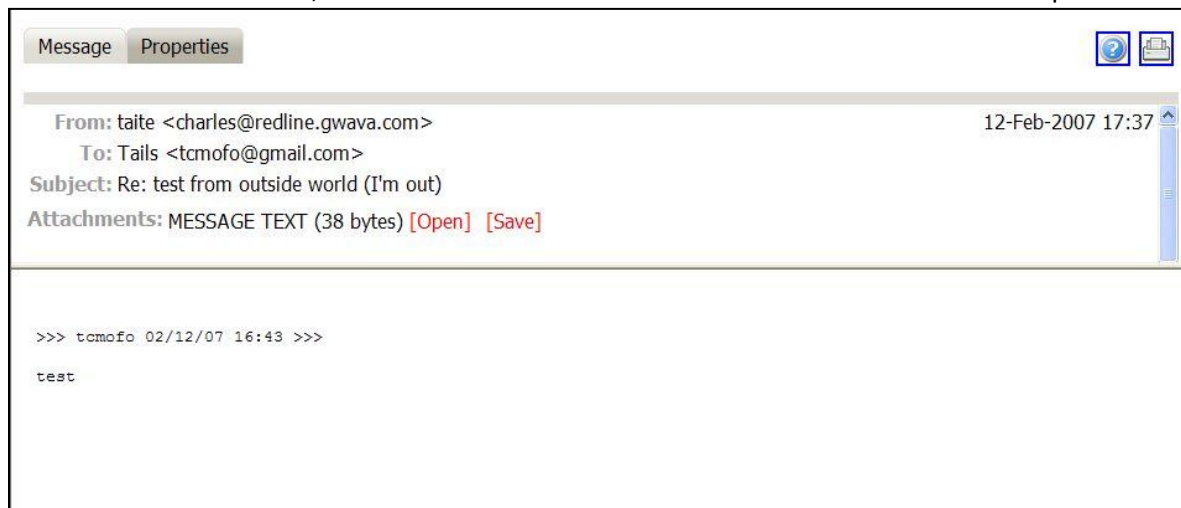
[See Search in More Detail](#)

Opening/Viewing a Message

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

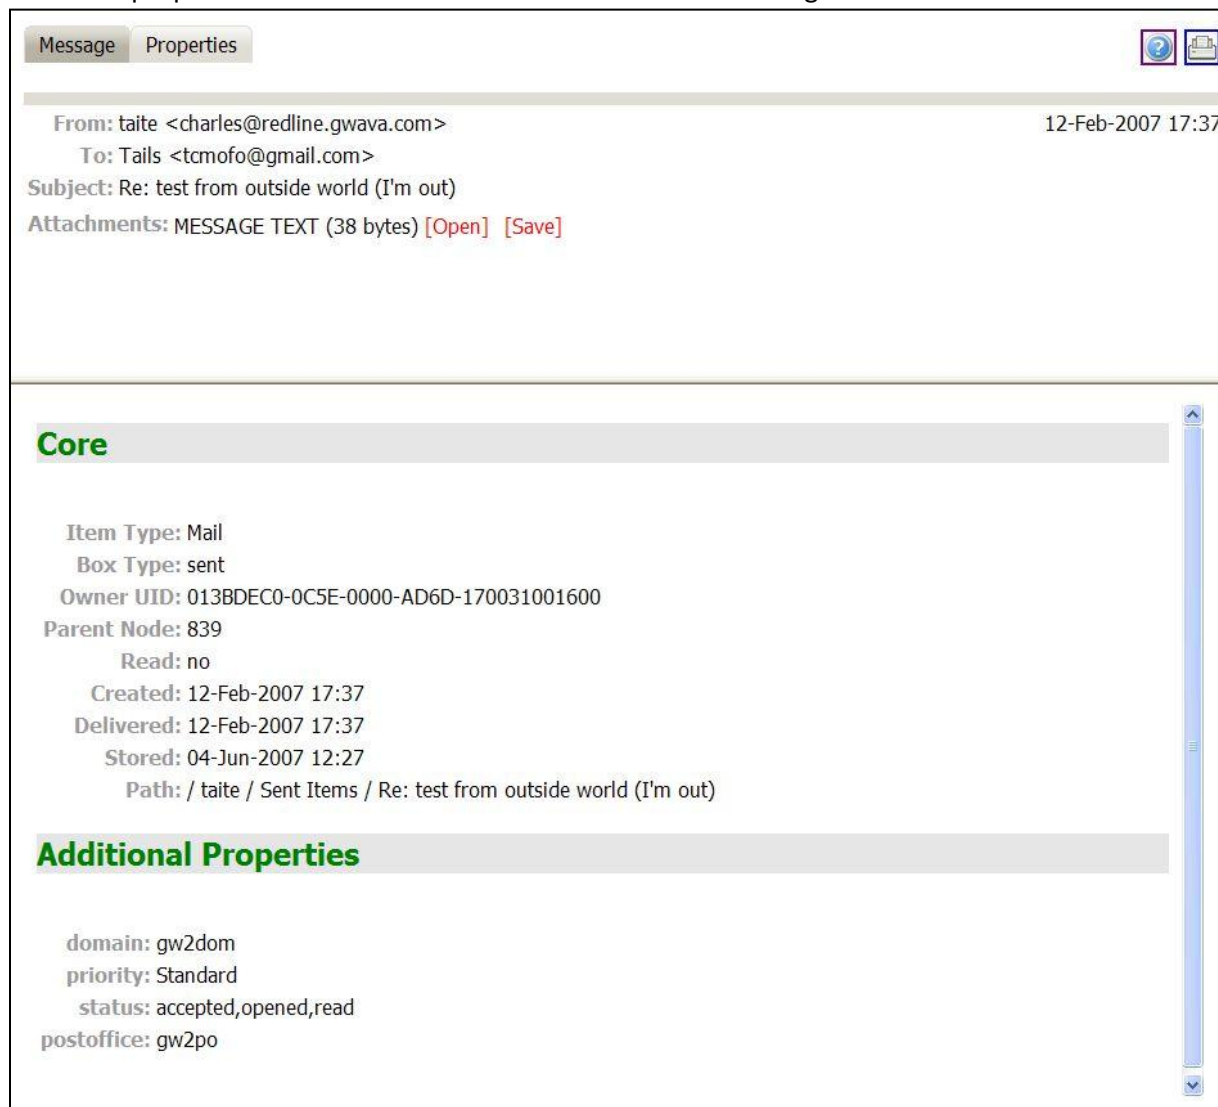


You will see all of the text, the headers and attachments as listed. Here is another example:



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:



The screenshot displays the 'Properties' tab of a message in the GWAVA Retain application. The interface is divided into two main sections: 'Core' and 'Additional Properties'.

Message Header:

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

Core Properties:

- 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 GroupWise 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.

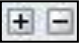
Core Tab:

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](#).

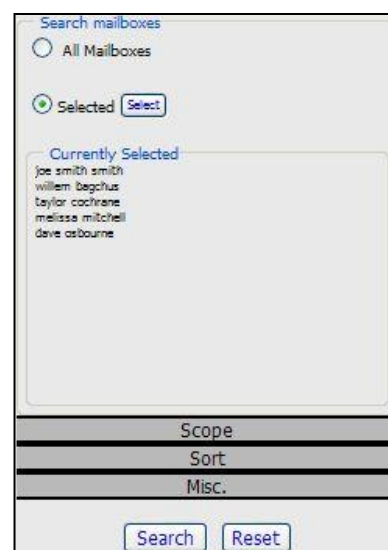
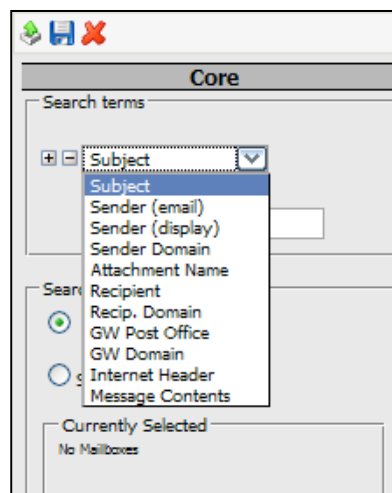
You can search based on all items shown on the right with operators including:

- 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 GroupWise 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 GroupWise 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 GroupWise, 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 GroupWise.

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 GroupWise. 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.

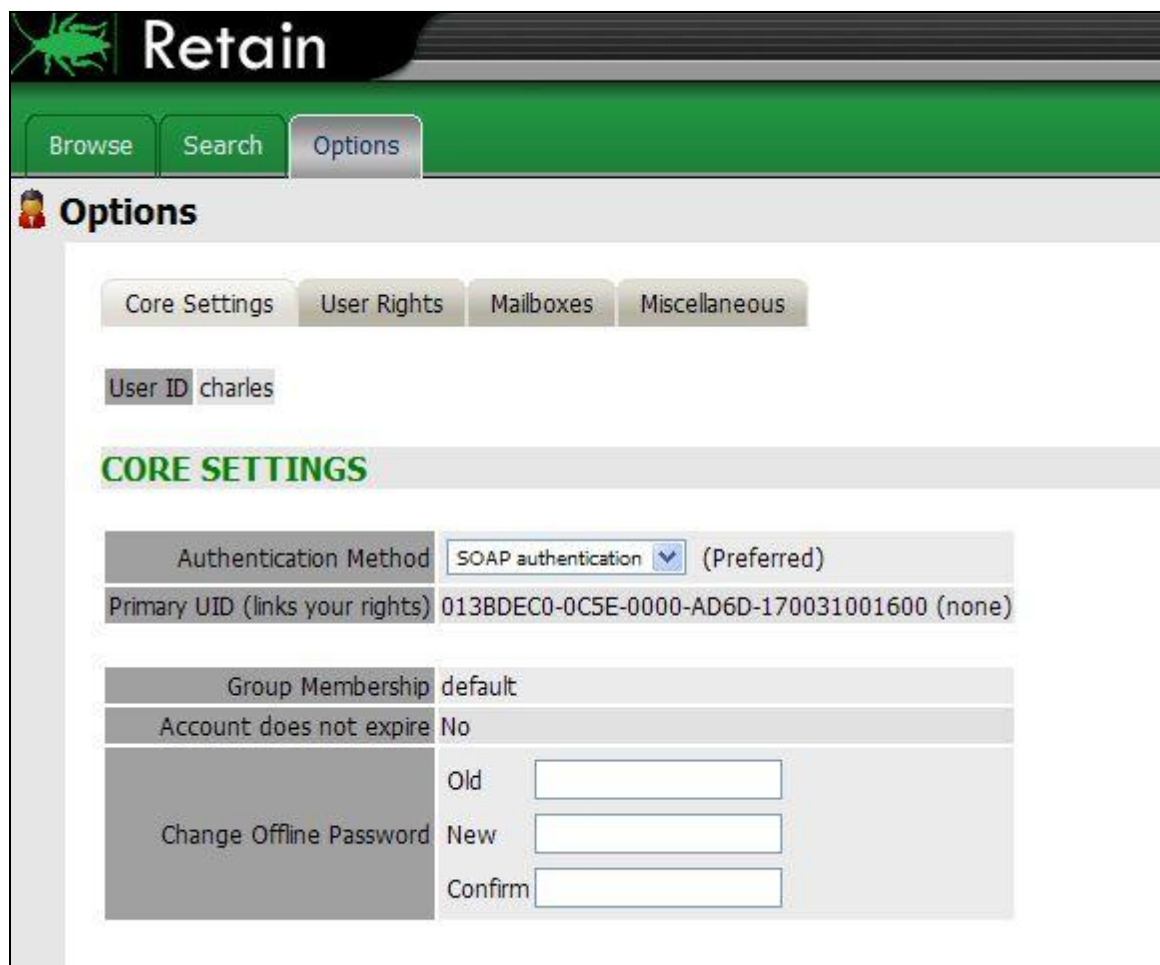
The screenshot shows the 'Sort' tab selected in a tabbed interface with 'Core', 'Scope', and 'Misc.'. Under the 'Sort Results' section, there are three dropdown menus: 'First By:' set to 'Creation Date', 'Then By:' set to 'Sender Domain', and 'Then By:' set to 'Recip. (display)'. At the bottom of the tab are 'Search' and 'Reset' buttons.

The screenshot shows the 'Misc.' tab selected in a tabbed interface with 'Core', 'Scope', and 'Sort'. The title '(Appointment, Task)' is displayed. Under the 'Start Date' section, there is a dropdown menu set to '(any date)' and a 'Range' button. Below that, under the 'End/Complete Date' section, there is another dropdown menu set to '(any date)' and a 'Range' button. At the bottom of the tab are 'Search' and 'Reset' buttons.

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 Retain application interface. At the top is a green header with the Retain logo and a navigation bar with 'Browse', 'Search', and 'Options' tabs. The 'Options' tab is active. Below the navigation bar is a sub-header 'Options' with a user icon. Underneath are four tabs: 'Core Settings', 'User Rights', 'Mailboxes', and 'Miscellaneous'. The 'Core Settings' tab is selected. Below the tabs, the 'User ID' is displayed as 'charles'. The 'CORE SETTINGS' section contains several fields: 'Authentication Method' is set to 'SOAP authentication' (Preferred); 'Primary UID (links your rights)' is '013BDEC0-0C5E-0000-AD6D-170031001600 (none)'; 'Group Membership' is 'default'; 'Account does not expire' is 'No'; and 'Change Offline Password' has three input fields for 'Old', 'New', and 'Confirm'.

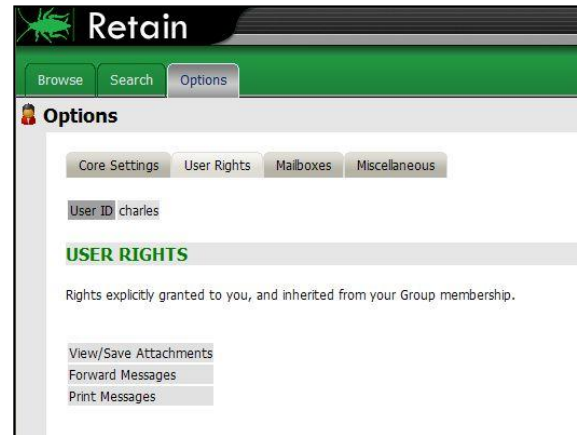
Core Settings

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

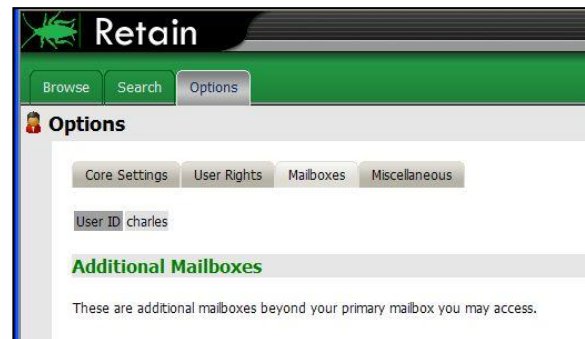
User Rights Tab

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.



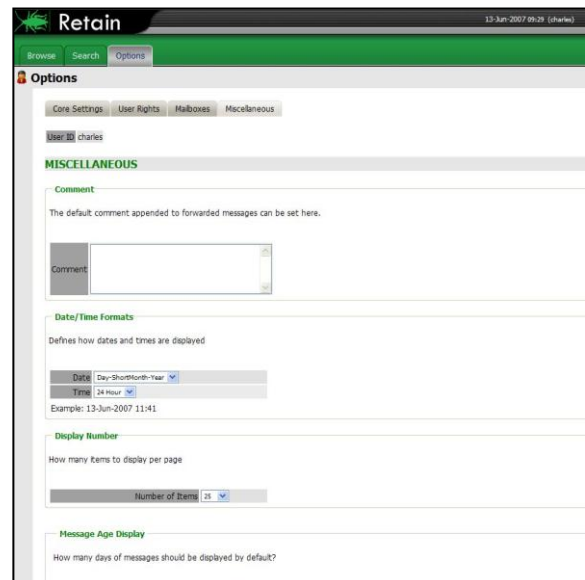
Miscellaneous Tab

This contains your display and session options.

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

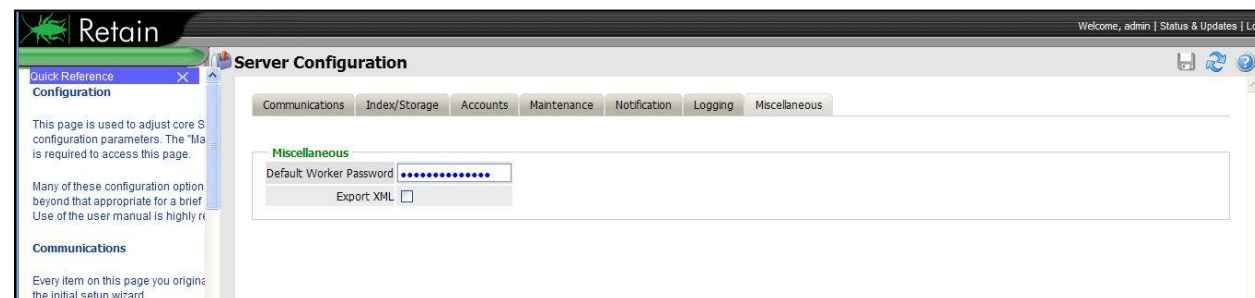
Options are:

- **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.

Credits 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 bug logo in the top left corner.



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	1d : 17h : 19m
currentTime	Thu Oct 18 08:29:49 EDT 2007
timezone	Eastern Standard Time
retainName	Retain-retaindemo
jvmVersion	1.4.2_11-b06
jreVersion	1.4.2_11-b06
javaHome	/usr/lib/jvm/java-1_4_2-sun-1.4.2.11/jre
tomcatHome	/srv/www/tomcat5/base
logHome	/srv/www/tomcat5/base/logs
osArchitecture	i386
osName	Linux
userName	tomcat
tmpDir	/var/cache/tomcat5/base/temp
tomcatVersion	Apache Tomcat/5.0
tomcatFreeMemory	10944
tomcatMaxMemory	65088
tomcatTotalMemory	47584
retainHome	/opt/beginfinite/retain/RetainServer
retainEdition	Standard
retainBuild	46
retainAvailability	beta
derbySchema	2
messageStoreSchema	1
indexingSchema	1
derbyVersion	10.2.2.0
luceneVersion	2.2.0
hibernateVersion	3.2.2.ga

Core Info

JAR List

Jobs

Schedules

Profiles

Workers

Quartz

Retention Date Utility

Appendix A:

Backing up Retain

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

- ASConfig.cfg
- Derby database (Embedded database)
- 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 index and embedded database backup under [Maintenance](#). Retain will automatically backup ASConfig.cfg, the Derby database, and the Indices. 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

To backup the SQL database is dependent on the SQL database you chose to use. 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 redaim 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 (5.5.2.6), which ships with Retain, Edit the following configuration file:
`/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 RetainServer, and 256MB for the RetainWorker. 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.

NOTE: Windows can allocate a maximum of 2 GB of memory to Java and Tomcat.

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.*

Every OS has a different methodology! Please read carefully.

Netware/OES NetWare:

- 1) Make a backup copy of **tomcat4.ncf**.
- 2) Edit **tomcat4.ncf**. You'll find these command line switches, that will look something like this:
 - -Xmx256m 256MB maximum memory by default. Change to 1024MB, for example.
 - -Xss160k 160KB stack size
 - -Xms512m 512 MB initial memory allocated.

SLES 9:

- 1) Edit **/usr/share/tomcat/catalina.sh**.
- 2) Add a line right below the **#!/bin/sh** like (for a Retain Server):
CATALINA_OPTS="-Xms512m -Xmx1024m -Xss160k"

OES Linux:

If you are using the OES Tomcat4 instance (default)

- 1) Edit **/etc/opt/novell/tomcat4/tomcat4.conf**.
- 2) Edit **JAVA_OPTS** (may contain some settings already) to something like: (for a Retain Server)
JAVA_OPTS="-Xms512m -Xmx1024m -Xss160k"

***** A typical OES Linux server uses lots of web services and if you installed them, don't skimp on RAM.**

If you are using the SLES Tomcat5 instance, apply the SLES 9 instructions.

SLES 10:

- 1) Edit **/etc/sysconfig/j2ee**
- 2) Change the **CATALINA_OPTS** line to something like: (for a Retain Server)
CATALINA_OPTS="-Xms512m -Xmx1024m -Xss160k"

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 Some Common Error Messages

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
- Disabled/Expired account

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.

It contains:




The 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 migration from a single workstation, and not involve the users at all.

See [Appendix D](#), below for more information.




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.

Diagnostic Tester

The Diagnostic Tester is a tool for troubleshooting and diagnostics. You can copy the code to your retain server directory, and you will be able to go to a user's live mailbox, read mail items, and test archiving. Consult the readme file for further information on its use.

Recover Admin

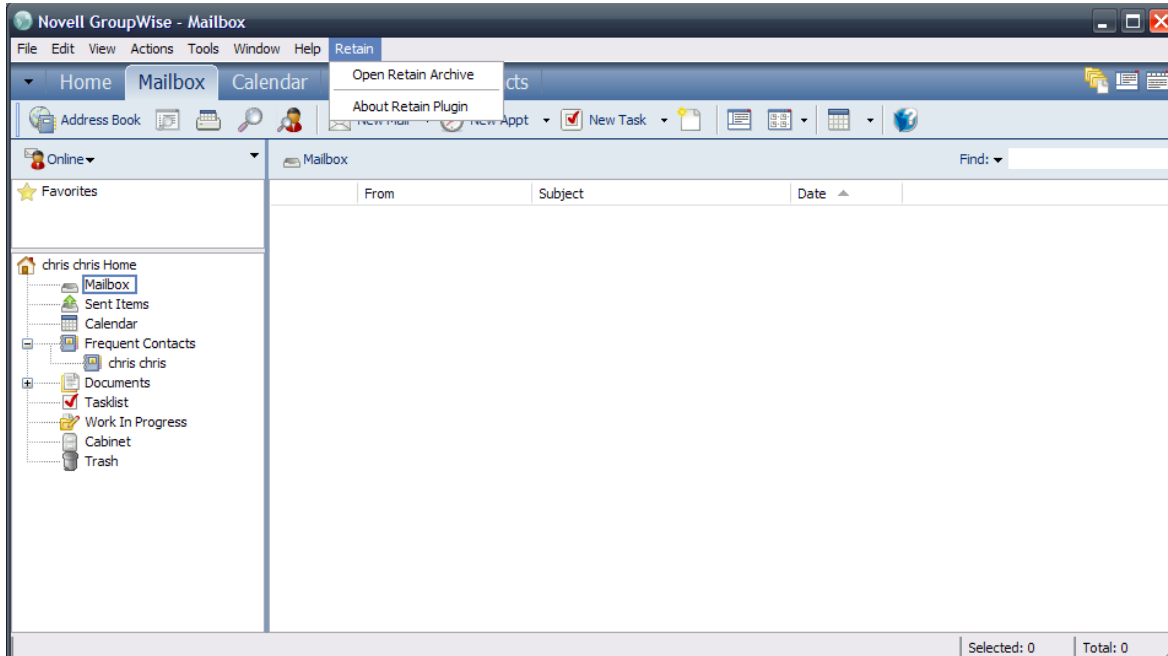
This is an emergency tool you can use if you happen to lose your admin account/password.

This utility will let you create a new admin.

Consult the readme file for further information on its use.

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. (.../<Retain180>/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
- (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 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 -jar retain_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>

Remote systems **require** 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.

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)
```


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

Appendix D – 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.

Workers Save Changes

Worker: migrator1

[Edit Worker](#) [Remove Worker](#)

No jobs are associated with this worker.

Enter new worker below:

Worker:

[Polling](#) [Logging](#) [Connection](#)

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.

Server Protocol:

Server Host Name:

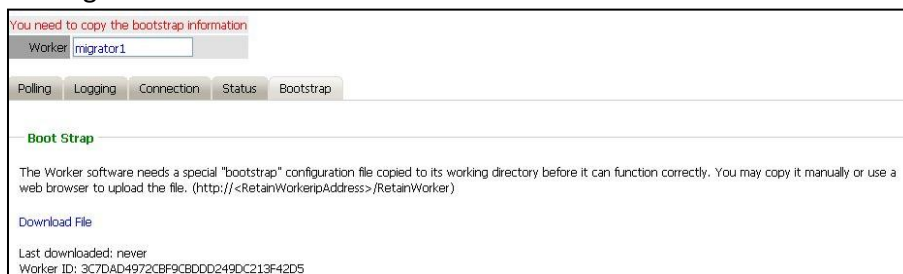
Server Port:

Server Path:

Worker Password:

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

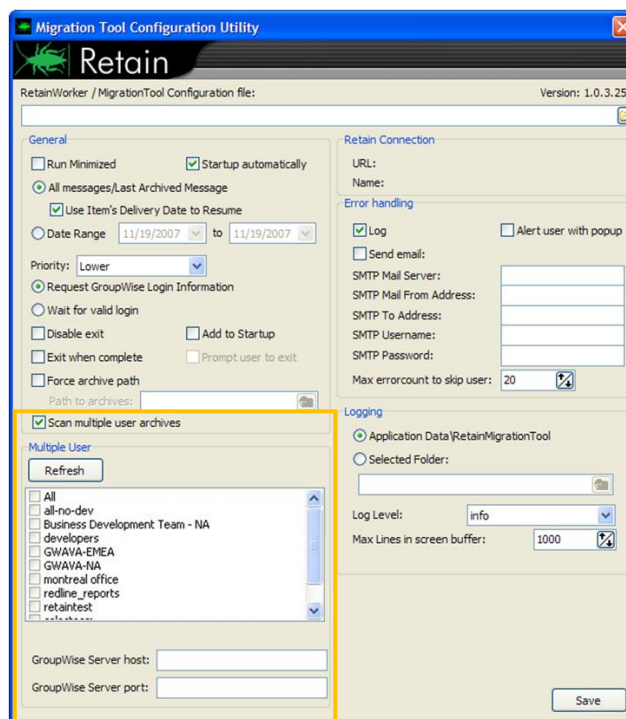
- 4) Once you've saved changes, you will be able to save the bootstrap file. Save it in the directory of the migration tool.



- 5) Now start up the migration configuration utility. At the very top of the screen you will see a drop-down selector to tell it where to find the bootstrap file you have just saved.

General Options

- Run minimized. Sometimes, you might want to run the migration in the background while the user is doing other things.
- Startup automatically. Select this to make the migration tool start the migration process as soon as it is loaded.
- The Force Archive Path option allows you to direct the archive migration utility to a directory where multiple archives have been copied. The migration tool will try to migrate all archives found in this specified directory.



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.

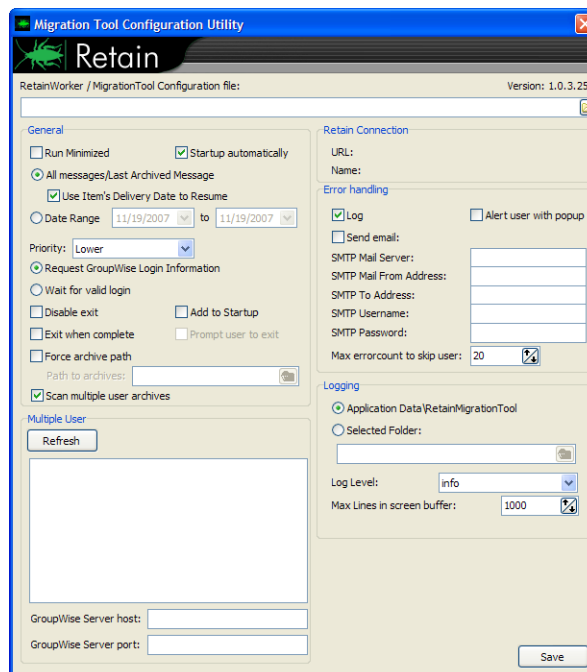
- Use item's delivery date to resume. If the migration process is interrupted and this option is selected, the migration tool will pick up where it left off the next time it runs.
- Request GroupWise Login Information. If you run the migration tool when the GroupWise client is not logged in, it will start up GroupWise and prompt you for login credentials.
- Wait for valid login. If you run the migration tool when the GroupWise client is not logged in, it will wait for the user to log in to GroupWise and then it will begin archive migration.
- Disable exit. This will prevent the user from closing the migration tool while it is running.
- Add to startup. This will add the migration tool to the startup so it begins when the workstation is booted. This is useful for long migrations where you want the tool to continually pick up where it left off and keep running while the user is logged in.
- Exit when complete. This option will force the migration tool to quit once the migration is complete. This is useful for unattended long migrations. You can also prompt the user before exiting.

Retain Connection.

- a) This information is a confirmation of what is in the bootstrap file you created for this migration. It lists the URL that the migration tool will use to talk to the Retain Server.

Error Handling

- a) Log. Choose whether you want the migration session to be logged.
- b) Alert user with popup. In case of an error, choose this option if you want the user to be notified of an error.
- c) Send email: If you want the user to be notified of an error by e-mail, choose this option and then fill in the SMTP details so it can send the message. This can be your GWIA if you wish. Be sure you can relay through this SMTP server!
- d) **IMPORTANT! Max Error Count to skip user:** The Archive Migration will skip the current user after 20 errors, unless you specify a higher number. *Re-migrating a user's archive will result in duplicate messages in the Retain Database.* To skip all messages with errors and continue archiving all available mail for each user, you must set the error number to '0'.



Logging

- a) Choose where to store the log file.
- b) Chose the log level, what severity of message would you like to appear in the log?
- c) Max Lines in screen buffer. Choose how many log lines will be in the scroll buffer that will appear on screen. This does not affect the size of the log FILE.

Multiple Users

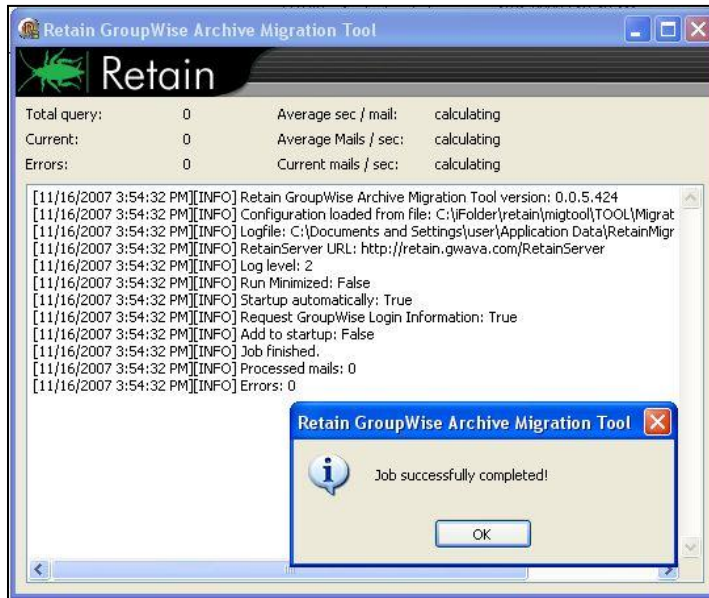
- a) Select "Scan multiple user archives" to enable multi-user mode.
- b) Refresh. If you click refresh, it will ask the GroupWise client to return the distribution lists in the GroupWise system. The server used is by default the one the GroupWise client is currently logged into.
- c) You may specify another GroupWise POA to log into.
- 6) Finally, click SAVE to save the migration tool configuration file. Save it into the directory where you will run the migration tool itself.
- 7) Run the migration tool when ready. It will read the configuration file and it will attempt to start the migration.



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. This is why it is critically important that the workstation you're using 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.

Appendix E – 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 Retain 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.

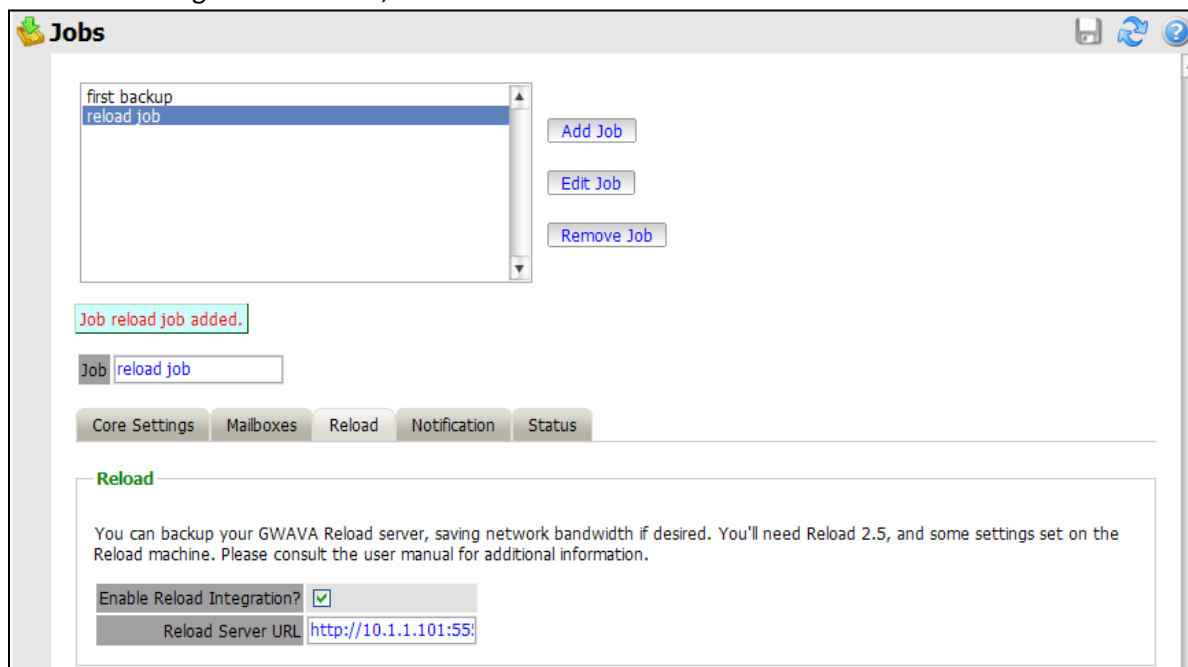
Timing

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.

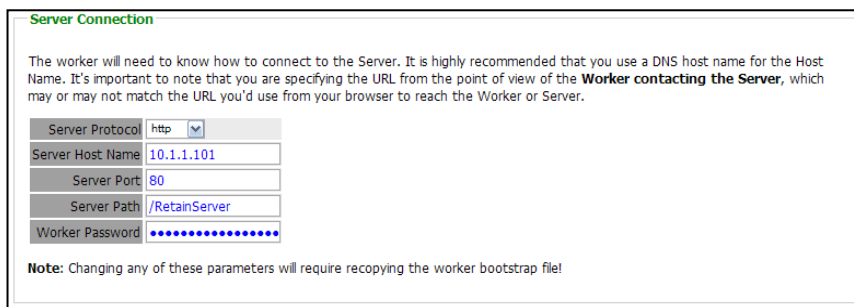
Enter the management console, and select Jobs from the Data Collection menu.



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.

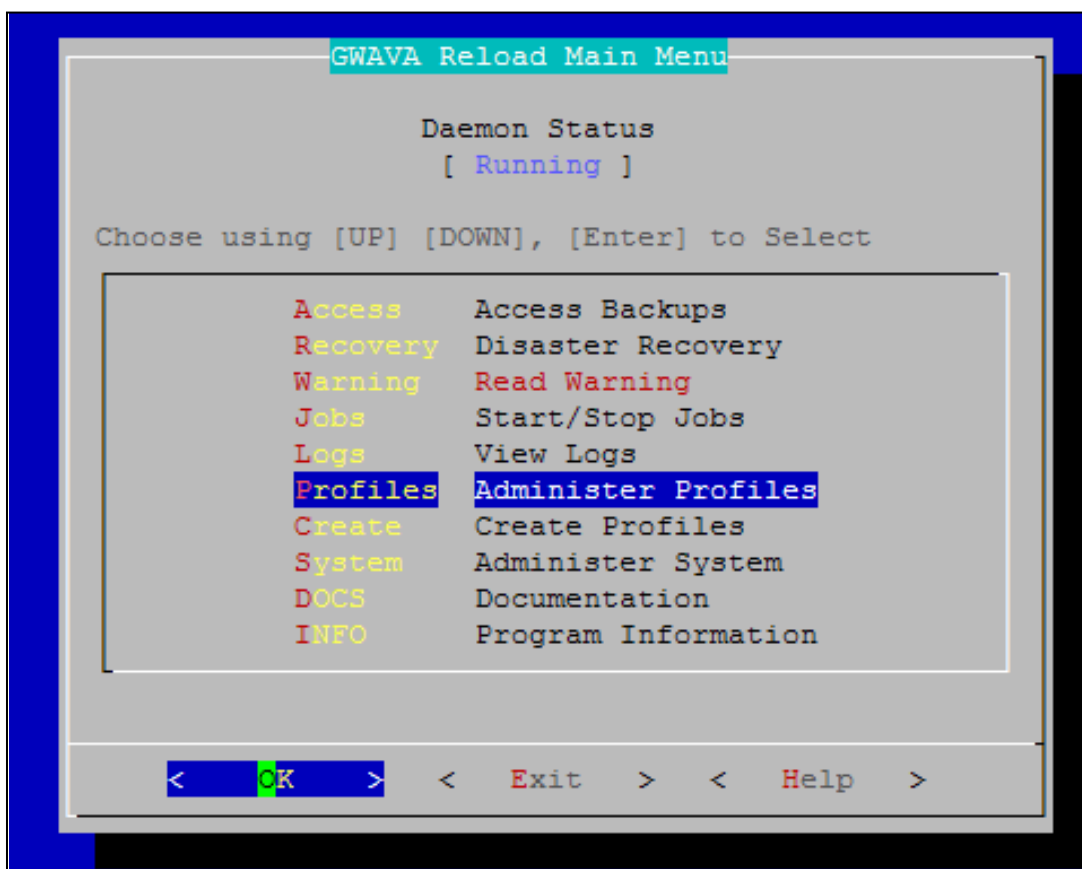
First, Reload must be set up so that the backups are available. There is a special feature in Reload for this. It calls up a post office agent that stays up all the time and it only goes down long enough to change to the most recent backup. So it will always be there with very small interruptions as 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.

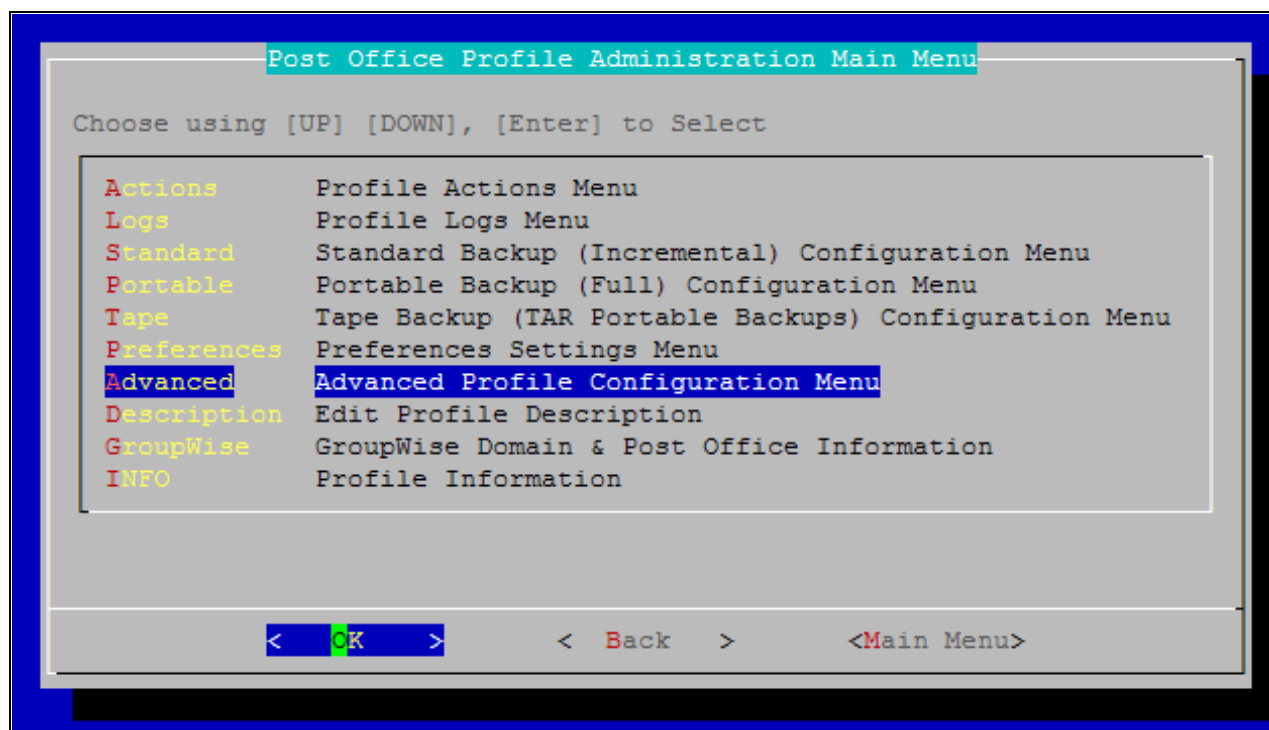
- 1) Start up Reload's Administration menu.



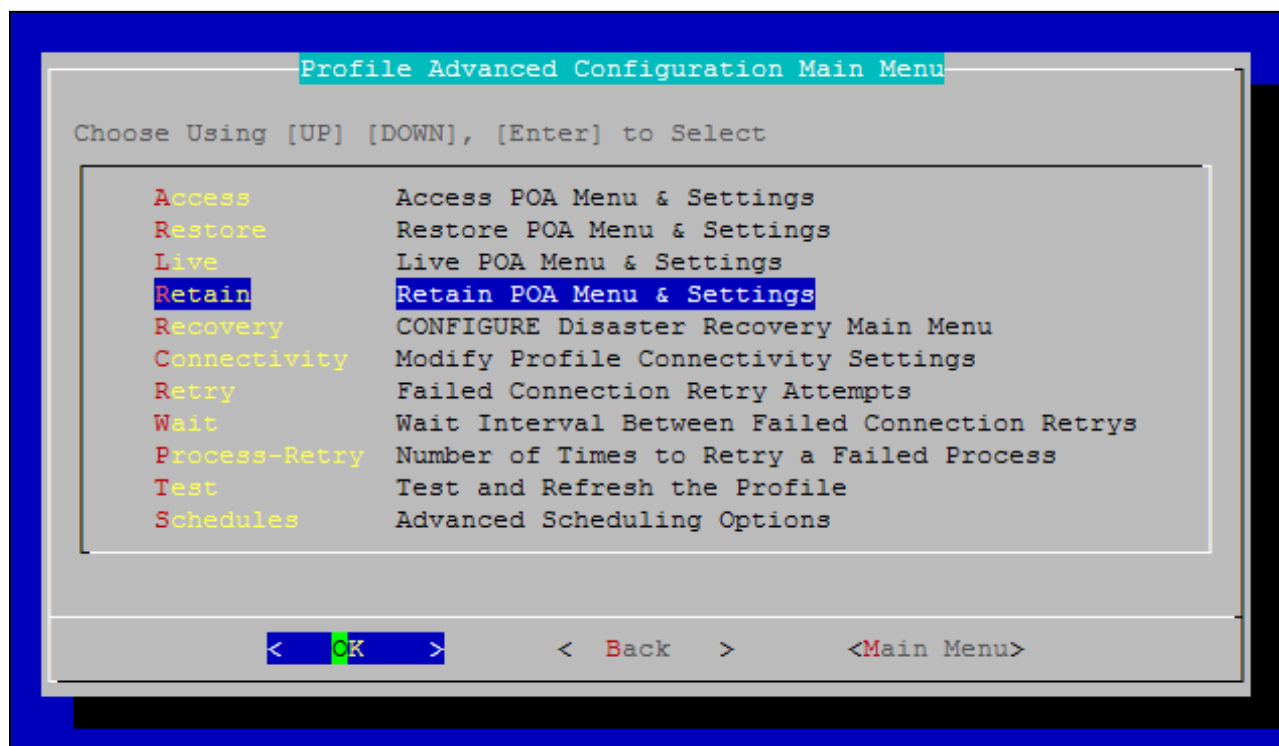
- 2) From the main menu, choose Profiles – Administer Profiles.

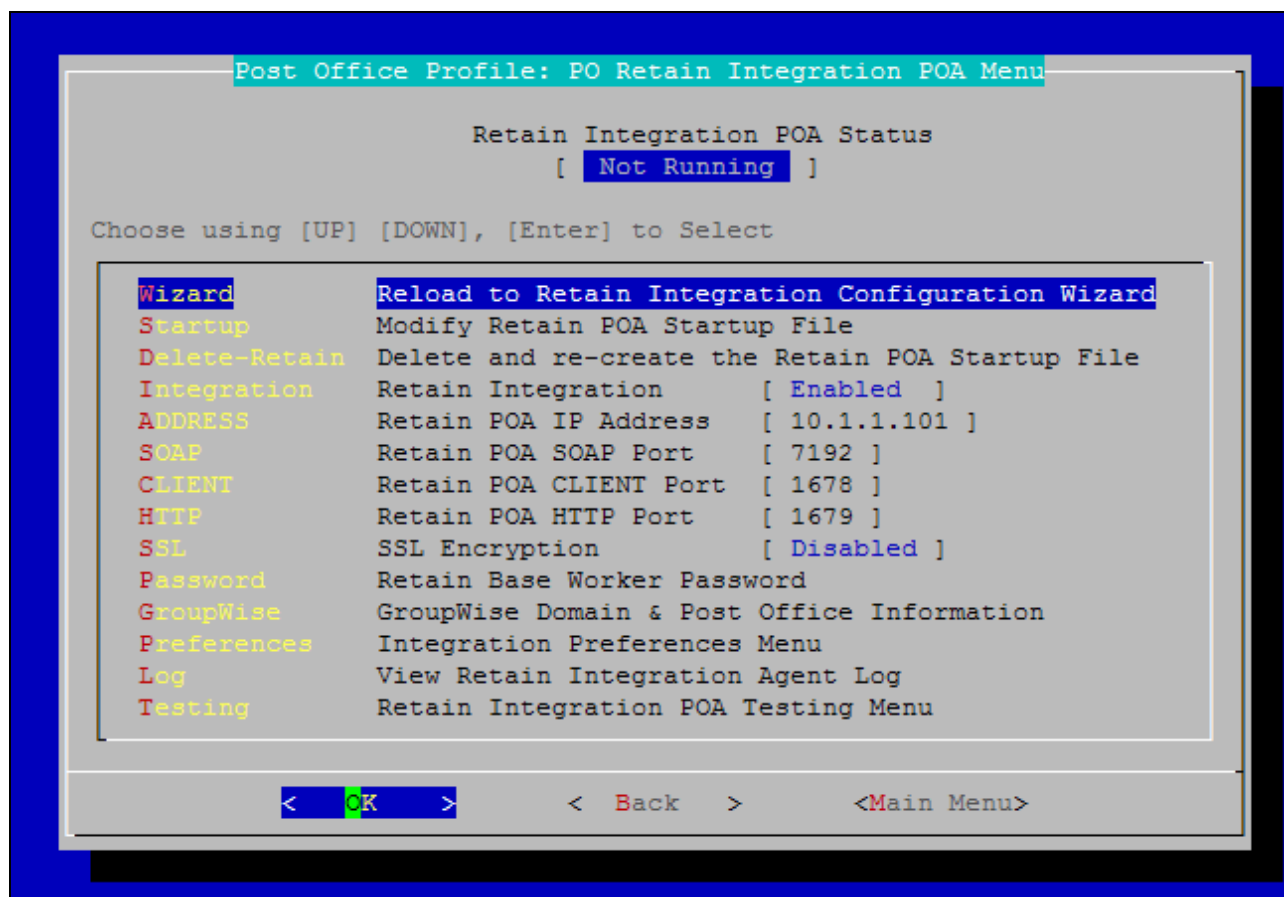


- 3) Choose Advances Profile Configuration Menu



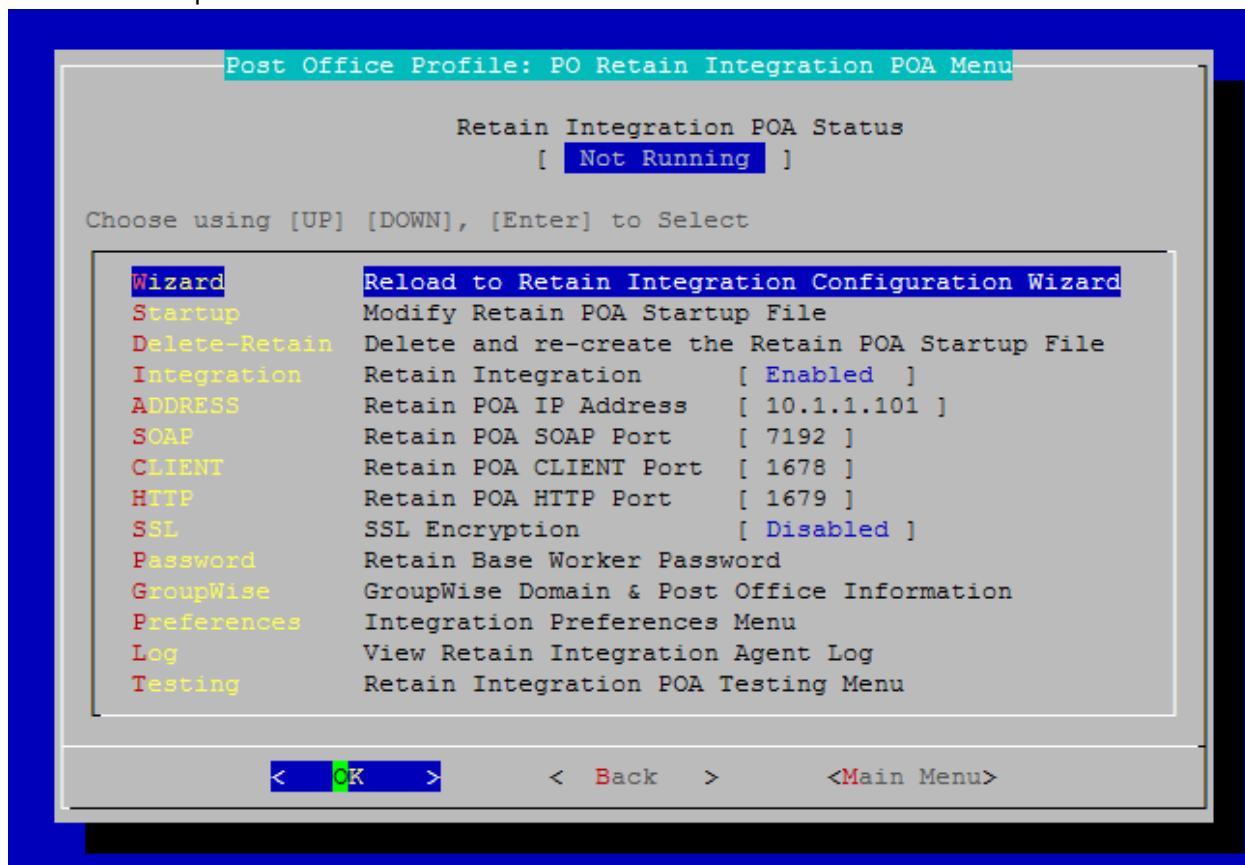
4) Choose Retain POA Menu & Settings





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 dialog boxes. The top dialog box is titled "Retain POA SOAP Port" and 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 buttons labeled "< OK >" and "<Cancel>". The bottom dialog box is titled "Retain POA CLIENT Port" and 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 buttons labeled "< 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.

The screenshot shows the 'Profiles' window in GWAVA Retain. A list on the left contains 'default' and 'reload', with 'reload' selected. To the right are buttons for 'Add Profile', 'Edit Profile', and 'Remove Profile'. Below the list, a message states 'No jobs are associated with this Profile.' A status bar indicates 'Profile reload saved.' The 'Profile Name' field contains 'reload'. Below this are tabs for 'Core Settings', 'Scope', 'Miscellaneous', and 'Advanced'. The 'Scope' tab is active, showing the 'Date Range to Scan' set to 'All Messages (ignore date)' and the 'Duplicate Check' section. The 'Duplicate Check' section explains that Retain stores a single instance of each message and attachment, and defines how the Worker determines new items. It includes a dropdown menu set to '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.

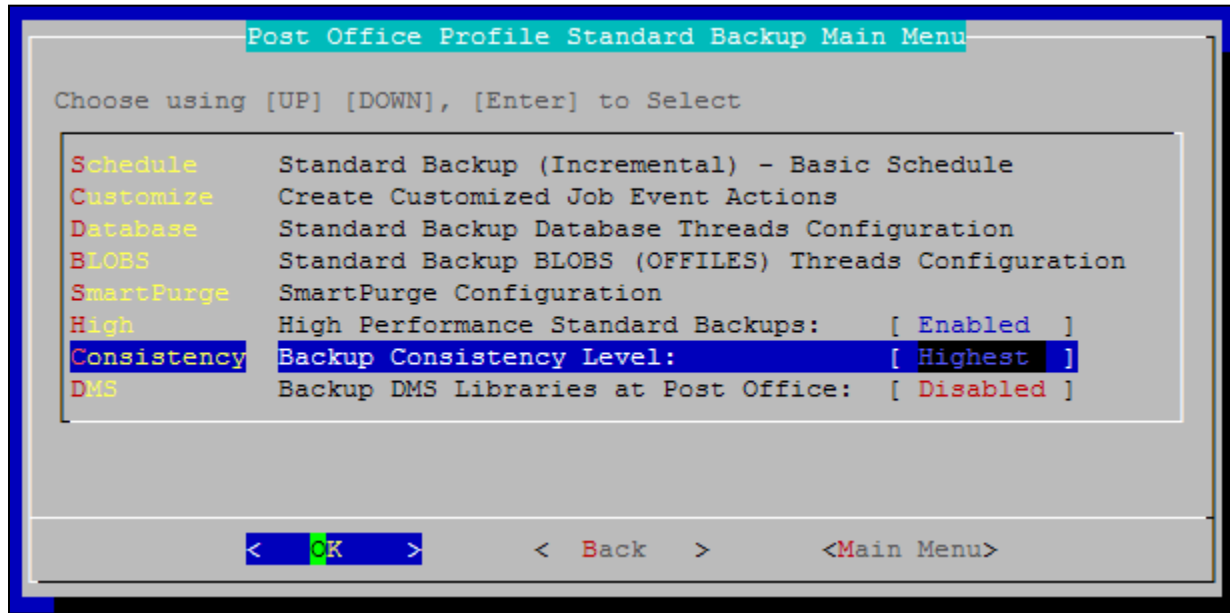
The screenshot shows the 'Set Storage Flags' section. It contains a paragraph explaining the function of the flags: 'If you are using either the Purge or Retention features in GroupWise, you probably want these to be advanced automatically as items are stored, so users may delete messages in their mailbox 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.' Below the text are three checkboxes: 'Retention Flag' (unchecked), 'Purge Flag' (unchecked), and 'Item Store Flag' (checked).

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 F – 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+rwX <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:

```
<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 the RetainServer 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.