

Overview

What Retain does

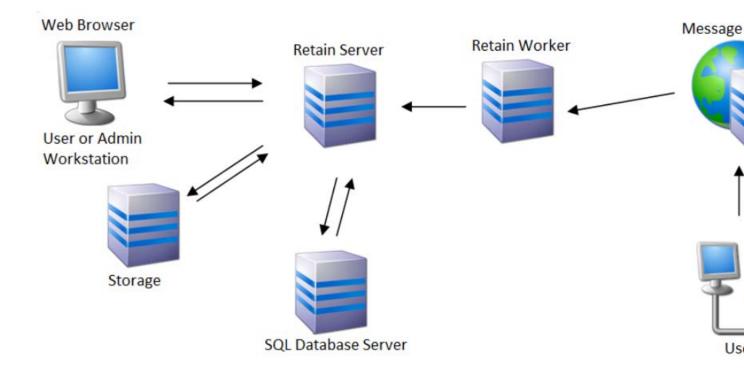
Retain provides a service of long-term storage of data as well as providing search, retrieval, and review services for archived messages. Retain is NOT a backup or emergency restoration system. Retain archives messages and data from messaging systems, phones, and social websites and stores the data for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

How Retain works

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

Architecture

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



- Retain Server: This is the core piece of Retain. All functions are controlled from the Retain Server.
 The server also manages the Retain Worker and stores data in the database server. There is only one Retain Server per system.
- **Retain Worker:** There is at least one per system, and often one per messaging server. The Retain Worker performs the data collection and transfers the collected data is handed to the server. The server stores the data in a database.
- **SQL Database Server**: This is an SQL server where Retain stores the collected data. It is not actually part of Retain. Retain was designed to support many different databases. NOTE: Installation, maintenance, tuning, and backup of the database is the customer's responsibility.
- **Stubbing Server:** The stubbing server works to remove large messages from the GroupWise system and creates a 'stub', or link, to the message which is stored in the SQL database. See the Stubbing Server section in the Main Guide to decide if a stubbing server is correct for your system. From a User's point of view, there is no change to the behavior of their GroupWise mailbox. Currently stubbing is only supported for GroupWise 8.01 or later
- Retain Router: The Retain Router gathers messages data from Android and BlackBerry phones using the Retain application, and is installed and controlled by the local administrator. Phone data is sent to GWAVA Inc. servers which check the specific identifying number connected to each phone, and forwards the data to the Retain Router to which the device ID belongs. The data is then stored until the Retain Worker retrieves the data and archives it into the Retain Server. Afterwards, the data is accessed the same as all other message data in the Retain Server.

Best Practices for Component Placement

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

How Retain Stores the Archives

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

Other components that Retain depends on

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

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

Design Considerations

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

SQL Database Server

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

- Network link between the Retain Server and SQL Database Server must be speedy and reliable.
- The Retain database may have to be manually created by the administrator and a user account must be assigned with full rights. See the Database section.
- Storage requirements: Roughly equivalent to the cumulative size of the message data store. See the 'Estimating Storage Requirements' section.
- NOTE: Installation, maintenance, tuning, and backup of the database is the customer's responsibility. Tuning an SQL Database Server can result in significant performance gains.

Retain Server

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

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

Retain Worker

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

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

Retain Router

The Retain Router needs to be visible and accessible to the Internet in order to receive phone data.

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

Messaging System Address Book

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

Retain Services

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

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

Email Systems

- Exchange
- Office 365
- GroupWise
- Gmail

Smart Phone Communications

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

Social Platforms

- Facebook
- Twitter
- Yammer
- Linkedin
- Yahoo Messenger

- Google Talk
- Lync & Skype for Business
- AOL Instant Messenger

Supported Search Portals

- Google
- YouTube
- Yahoo
- Bing
- Wikipedia

Copyright © GWAVA Inc. 2015 All rights reserved.

"GWAVA" is a registered trade mark of GWAVA Inc, 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1
Retain" is a trade mark of GWAVA Inc, 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1
GroupWise is a trademark of Novell, Inc.

Exchange and Windows are trademarks of Microsoft Inc. Facebook, Linkedin, Yahoo, AIM, and Twitter are registered trademarks of their respective owners.

Upgrading to Retain 4

To upgrade to Retain 4 requires that Retain be upgraded to 3.5+ first.

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

For Linux:

Retain 2.x to 3.x http://support2.gwava.com/kb/?View=entry&EntryID=2185 Retain 3.x to 3.5 or 3.5.1 http://support.gwava.com/kb/?View=entry&EntryID=2317

For Windows:

Retain 2.x to 3.x: http://support2.gwava.com/kb/?View=entry&EntryID=2226 Retain 3.x to 3.5.1 http://support.gwava.com/kb/?View=entry&EntryID=2433

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

Upgrade changes

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

Requirements:

Storage Space

While the migration is performed, a new index will be created, so the current index size will be doubled temporarily while two indexes exist. Once migration is complete, the old index may be removed and space reclaimed. Also, the new index may be larger than the current index, due to the increased power and abilities of the new indexer. How much change there is, depends on the composition of the current archive, any limits set in indexing attachments, and how many attachments there are in the archive. If there are a lot of attachments, or very large attachments, then the increase in size will be significantly larger than if there are small or a limited number of attachments in the archive. With no limits on indexing attachments and a lot of large attachments in the system, an increase of up to 4 times the current index size is possible. With limits included, the new index may be smaller after upgrading. For the upgrade and a worst-case scenario, ensure that up to 5 times the current index size is available before starting the index migration. (With a current index size of 5 GB, ensure that the volume has a minimum of 20 GB free before starting the migration.) The average increase is expected to be around 20%.

If there is insufficient space to perform the migration, Retain will be placed into maintenance mode until more space is provided.

Time and Access

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

System Requirements Changes

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

Licenses

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

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

Run the installer.

Select the second option: New Install/Upgrade option.

Select components to be upgraded or installed. Select all components desired. (Only components selected here will be active on the local system.)

```
Retain Installation Options

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

As soon as the Retain Message Router is selected or refused, the installation will begin. Wait until the installation is finished.

Browse to Server Manager



Database Migration

The configuration database schema has been enhanced in this version of Retain and will be unavailable for use for a short time while it is being updated. Please login as the administrator to continue.

Enter credentials

The upgrade requires a configuration database update.

Select the 'Enter Credentials' button and enter the administrator account username and password. Select the button to start the migration.



The database schema update is in progress and should not be interrupted. This message will be updated upon completion.

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

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



Database schema update complete. Please check your logs for error messages. If none are present, please restart Tomcat

Once the update has completed, restart tomcat to initiate the new database.

The tomcat restart will take a little bit longer than a normal restart, as the database needs to be initialized and set up the new interface.

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



Login to the Retain Management interface. New Licenses for Retain 4 must be immediately applied. These licenses should be provided by the sales representative. (Removing old licenses is not required, but may help to keep interface looking clean.)







Licensing Summary

General Summary

Fri Nov 13 22:55:01 MST 2015 Date Product GWAVA Retain 4.0.0.0

Total Active License Count 0 Total Inactive License Count 2424 Merged Mailboxes 10 -> 5

Individual License Breakdown

Active License Count Inactive License Count System

qw 0 2406 exchange 0 18

Licenses -

CN=pg-2358, EMAILADDRESS=info@gwava.com, O=GWAVA, L=Montreal, ST=N/A, C=CA X Licensed to

License ID 87d1b648-47c0-4c31-ac1e-18381d059b6c

License State Valid License file installed.

License Type gw

Tue Sep 29 12:16:59 MDT 2015 Install Date License expires Thu Aug 16 01:22:23 MDT 2035 Support expires Wed Feb 28 23:59:59 MST 2024

Your license file is currently stored at: /var/opt/beginfinite/retain/license/rtn7036907066489667840.pem Make sure you keep a backup of this file.

X Licensed to

CN=pg-2350, EMAILADDRESS=info@gwava.com, O=GWAVA, L=Montreal, ST=N/A, C=CA

License ID f2567fe9-86e2-44d3-9788-ac77a0e73e78

License State Valid License file installed.

License Type base exchange gw socialmessaging mobile blackberry

Install Date Tue Sep 29 12:16:18 MDT 2015 License expires Wed Aug 15 09:34:43 MDT 2035 Support expires Sat Jan 27 23:59:59 MST 2024

Your license file is currently stored at: /var/opt/beginfinite/retain/license/rtn1143912542230642963.pem Make sure you keep a backup of this file.

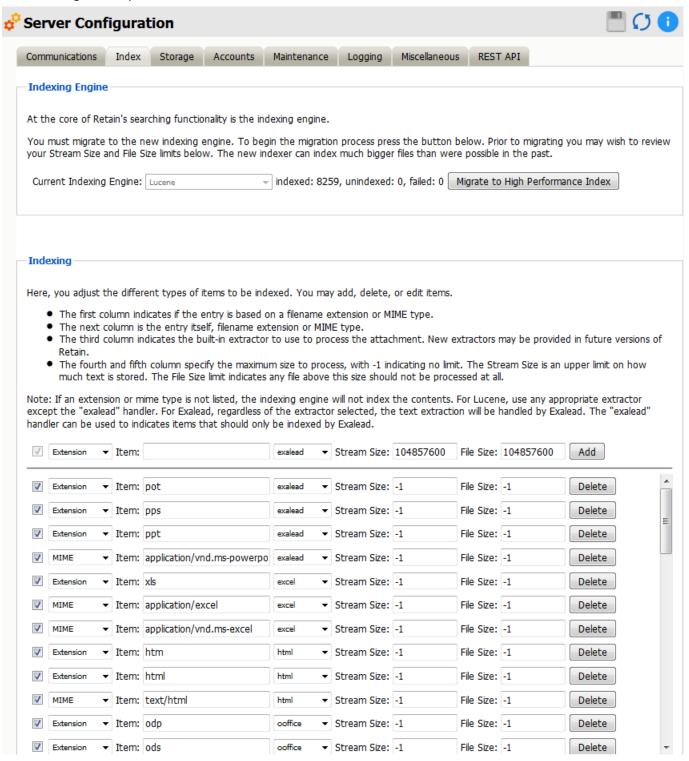
Submit License Browse... No file selected. Submit License

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

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

Once the licensing has completed, immediately head to the Server Configuration page and select the Index tab.

Server Configuration | Index



Select the Migrate to High Performance Index button.

Current Indexing Engine:	Lucene 🔻	indexed: 8259, unindexed: 0, failed: 0	Migrate to High Performance Index

There are two index options: High Performance Index, (or HPI), and the standard index. The default is set to the standard index. If you wish to use the HPI index engine, please contact support for guidance, system planning, and installation. The standard index engine is very powerful and sufficient for everything up to extremely large or busy systems.

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

Configure New Indexer				
Configure New Indexer				
You now need to specify an Inc brief summary of each of the ch		important choice, as you	r decision affe	ects the scalability, hardware requirements, and capabilities of Rel
on a single server it does High Performance Index be dedicated 64-bit serv	s not offer distributed processin er - This indexer is a set of 2 pr ers (quad core recommended)	ng power nor high availabil rocesses running on one o , with at least 16GBs of R/	ty. r more additio	s up when Retain starts and shuts down when Retain shuts dow onal servers. The 2 processes running on the servers of the high uire a separate installation using the Retain HPI installer and a HP
Would you like to upgrade to t	ne High Performance Indexer?	No, use Standard Indexer	•	
Port: 8081	(default usually available)			
Please enter Retain admin cred	entials:			
Username:		retain		
Password				

The Standard engine only requires the Retain admin credentials.

Enter the administrator credentials then save the changes by selecting the save-changes button to continue.



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

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

When the index page updates showing the new index is running, the new index has automatically become active and the upgrade is complete.

Technical Support

If you have a technical support question, please consult the GWAVA Technical Support section of our website at:

http://support.gwava.com.

Sales

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

100 Alexis Nihon Blvd. suite 500

Montreal, Quebec

Canada. H4M 2P1

Intended Audience

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

Organization

The guide is organized into two sections, Overview and Installation. The Overview section covers all the basic information needed to plan where to implement and install the different parts of Retain in your System. The Installation section briefly goes over the necessary steps to install Retain.

General use and management of Retain is covered in the User's Guide, please read that guide for setup and configuration steps and 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 and Exchange are copyrighted by Microsoft. GWAVA® is a registered trademark. Copyright © 2010 GWAVA Inc. All rights reserved. All other trademarks, service marks, and company names are properties of their respective owners.

Tomcat Embedded with Retain

Retain ships with embedded Tomcat. Retain requires you to install and use the embedded Tomcat. Installation on an existing Retain system running a standard Tomcat installation will remove the current Tomcat connection and install the embedded Tomcat that comes with Retain. Installation instructions are located in the install section.

Retain-Tomcat has changed the default memory tuning to 256MB in initial and 1024MB maximum memory, but the memory should still be tuned. See the Tomcat Memory Tuning section. The Retain install script, RetainInstall.sh, detects whether the embedded Retain-Tomcat installation is present or not, and, if missing, begins the embedded Tomcat install.

System Requirements

First decide where everything should be installed. Once the support and core components are installed, using Retain is the same no matter what platform it is installed on. You may spread out Retain's compo-

nents on different servers or run them all on one. VMWare ESX and ESXi is supported, but requirements are still in full effect and must be observed.

Retain may be configured in separate ways which heavily impact the requirements of the Retain server, indexing engine, database server, and locations dictate resource needs. Ultimately the amount of active users in the system determine the resource needs of the system.

Minimum System Requirements:

- Dual-Core 3 GHz or better Processor
- Server requires 12 GB RAM free
- Apache 2.2.x or IIS
- Supported SQL Database
- Java Development Kit (JDK) 1.8 64-bit
- As much storage space as messaging system requires.
 - NOTE: The Retain system requires a minimum of 10 GB of free disk space or it will enter maintenance mode. Messaging systems will have storage space requirements far exceeding the minimum and Administrators should look to the 'Estimating Storage Requirements' for minimum storage space estimation instructions.
- Java must be allocated a minimum of 4 GB of RAM.

NOTE: If the server has 8 or more cores, Java must be allocated a minimum of 5 GB of RAM. The remaining of the system requirements are determined according to system load as defined by the amount of active mailboxes and users being archived, (level of use per mailbox will vary). Amount of system utilization will determine the implementation of Retain.

Single Server

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

12 GB RAM

Remote Database

(Retain Server, Retain Worker, and Indexing engine on one server, separate Database server) 500-5,000 active mailboxes:

• 12-16 GB RAM

5.000-30.000 active mailboxes:

• 16-32 GB RAM

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

50,000 Active mailboxes

- 16+ GB RAM for Retain Server (depending on environment)
- Database on appropriate server
- Retain Worker is installed remotely on each post office in the message system

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

Supported Linux Versions

- Open Enterprise Server 2 Linux (64-bit)
- SUSE Linux Enterprise Server 12 (64-bit)

SUSE Linux Enterprise Server 11 (64-bit)
 NOTE: The Retain system requires more than simply space. Due to the storage of many small files, especially on systems archiving mobile message data, sufficient inodes are also an issue.

 SLES 11 changed the way the EXT3 file system defaults, and unless modifications are made, the EXT3 file system on SLES 11 will become insufficient for large storage systems. For information on the limitations and how to fix it, see http://support.gwava.com/kb/?View=entry&EntryID=2427. If a substitute file system is desired, the XFS file system is flexible enough to mitigate the risk.

Supported Windows Versions

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

Microsoft Exchange Requirements

- Exchange Server 2007 SP 3
- Exchange Server 2010 SP 2 or higher
- Exchange Server 2013

Novell GroupWise Requirements

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

Android Requirements

• 4.0.4+

Celltrust Secureline

• CellTrust Gateway 4.2.0 +

BlackBerry Requirements

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

Supported Social platforms

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

Supported Search

- Google
- YouTube

- Yahoo
- Bing
- Wikipedia

Supported SQL Database Servers

- MySQL 5.6x
- Oracle 10/11g
- MS SQL Server 2008 R2, 2012, 2014
- PostgreSQL 9

Supported Browsers

- Internet Explorer 10+
 NOTE: Retain's support of IE 10 will be limited after Microsoft ends support for IE10 on Jan. 12, 2016
- Firefox 34+
- Chrome 39+

Ports in use

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

Retain Server

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

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

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

Retain Worker

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

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

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

Retain Message Router

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

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

Other ports may need to be considered for NFS to work properly; see:

http://serverfault.com/questions/377170/which-ports-do-i-need-to-open-in-the-firewall-to-use-nfs. It is up to the customer to do this research in order to get NFS services to work properly.

Retain Reporting & Monitoring Server

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

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

Retain Stubbing Server

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

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

Retain for Social Media (RSM) Proxy Server

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

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

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

Database Management Systems

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

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

Blackberry

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

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

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

Other ports may need to be considered for NFS to work properly; see:

http://serverfault.com/questions/377170/which-ports-do-i-need-to-open-in-the-firewall-to-use-nfs. It is up to the customer to do this research in order to get NFS services to work properly.

Exchange

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

Google Apps Module

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

GroupWise

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

Mobile

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

Other ports may need to be considered for NFS to work properly; see:

http://serverfault.com/questions/377170/which-ports-do-i-need-to-open-in-the-firewall-to-use-nfs. It is up to the customer to do this research in order to get NFS services to work properly.

Office 365

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

CellTrust Secureline

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

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

Other ports may need to be considered for NFS to work properly; see:

http://serverfault.com/questions/377170/which-ports-do-i-need-to-open-in-the-firewall-to-use-nfs.

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

Estimating Storage Requirements

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

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

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

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

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

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

Tomcat Memory Tuning

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

Retain Server alone:

Minimum heap: 2 GBMaximum heap: 4 GB

Recommended Minimum system RAM: 8+ GB

A minimum of 8 GB of RAM is required. For larger systems, 16 GB RAM or more 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 other server processes are ideally minimal. Also, remember this memory is shared with other web applications.

Retain Worker Component(s)

Retain Worker alone:

Minimum heap: 2 GBMaximum heap: 3 GB

• Recommended Minimum system RAM: 4+ GB

The Retain Worker component(s) have been designed to be light weight. Nonetheless, a minimum of 2 GB RAM is recommended, and 3 GB 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.

Retain Server, Worker, and Database

Retain Server, Worker, Database all combined on one server:

Minimum heap: 4 GBMaximum heap: 6 GB

Recommended Minimum system RAM: 12+ GB

Combining the Retain system all on one machine is possible but not recommended for large systems. A minimum of 12 GB RAM is required, assuming that the Database system only requires 2 GB of RAM. System memory requirements will vary depending on the database system used, but may require 18+ GB of RAM. It is recommended for larger systems to maintain the Database on a separate server.

Tomcat Configuration Parameters

There are three significant parameters:

- 1. Memory allocated upon Tomcat startup
 - Indicates how much memory is immediately allocated and reserved to Tomcat upon startup.
 This memory will be in use for the entire lifetime of Tomcat and never available to the other server processes.

• Typically this is set to 50%-100% of the maximum memory parameter discusses below. It can be lower, but pre-allocating a sizable percentage of memory enhances performance and reduces memory fragmentation.

2. Maximum Memory available to Tomcat

- If the memory usage grows beyond the startup allocation, Tomcat will allocate additional blocks of memory in chunks as needed up to this limit. It will never return this memory to the general server memory pool. (Although memory internally will be freed and reused for Tomcat applications).
- Tomcat guarantees that it will never exceed this memory allocation parameter. If Tomcat runs out of memory, it will try to reclaim unused memory via garbage collection. If this is insufficient, the web application will be denied the memory allocation. Unpredictable (but invariably unpleasant) results will then occur.
- This is the most critical parameter to tune.

3. Stack Size

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

For Linux, edit the following configuration file:

/etc/opt/beginfinite/retain/tomcat7/j2ee

Retain-Tomcat is pre-setup to use minimum specs. To tune the system, edit the following lines with the appropriate parameters for your system:

```
CATALINA_OPTS=""
STACK_SIZE="-Xss256k"
MEMORY_HEAP="-Xms4g -Xmx4g"
CATALINA OPTS="$MEMORY HEAP $STACK SIZE"
```

By default, the embedded Retain-Tomcat is pre-tuned for basic functions. This will not be enough in a full production environment. Please tune to fit your system needs.

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

Checking memory:

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

SQL Database Guides

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

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

MySQL – WARNING: MySQL 4.x and 5 are provided with many SuSE Linux versions and are installable via YaST. Do NOT use them; they are NOT compatible with Retain and will corrupt your data if used. Uninstall MySQL 4 or 5 and get MySQL 5.6x. MySQL 5.6 needs to be downloaded and installed.

MySQL Quick Start Guide

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

To use MySQL with Retain, you will need to do the following:

- 1. Download MySQL
- 2. Install it to the server of your choice
- 3. Log in and create a database for Retain to use
- 4. Create a user with rights for Retain
- 5. Retain will do all the rest

Getting MySQL

Although SLES comes with MySQL 5.x as an optional package to be installed by YAST, the version supplied is out of date and prone to data corruption. It is strongly recommended you download and install the latest RPMs from MySQL. For MySQL to be useful, three packages need to be installed: the Server, Client, and shared client libraries.

- Go to http://www.mysql.com
- Get MySQL 5.6

Get the binaries appropriate for your system. Again, the Server package alone will not create a full working system; the Client and shared Client Libraries are required.

NOTE: 5.6 is the currently tested and supported version of MySQL, other versions may work, but are not supported by Retain.

A complete download and installer script for MySQL is provided for Linux. If specific instructions are required for the installation of MySQL on Windows, please see the MySQL site for proper instructions. https://dev.mysql.com/doc/refman/5.6/en/mysql-installer-gui.html

MySQL Optimization

MySQL requires careful memory tuning. By default it assumes it is running on a 64 MB Server. The administrator should give at least 1 GB RAM to MySQL. Larger sites may well allocate 2, 4, 16, or even 64 GBs of RAM depending on their data store. We do not support or discuss memory tuning requirements. How-

ever the following links may help (Retain uses InnoDB tables):

- http://www.mysqlperformanceblog.com/files/presentations/UC2007-Innodb- <u>Performance-Optimization.pdf</u> covers InnoDB specifically. Clear and concise instructions for optimization: http://www.mysqlperformanceblog.com/files/presentations/UC2007-MySQL-Server-Settings-Tuning.pdf
- Also, for basic optimization directly after MySQL installation, see: http://www.mysqlperformanceblog.com/2006/09/29/what-to-tune-in-mysql-server-after-installation/

MySQL Installation Via the Installer Script

Retain comes with automated MySQL installation and database creation scripts, mysqlinstall.sh and mysqlsetup.sh. The mysqlinstall script can download and install MySQL, or can automatically install if the files have already been downloaded. To use only the setup script, both the MySQL server and client packages must already be installed. The scripts are run during Retain Install *if* the MySQL database server is selected *and* the MySQL server will be housed on the same server as Retain. If the install initiates the scripts, they will be run back to back and do not need to be initiated manually. To utilize the scripts on a separate machine from the Retain Server, simply copy the Retain source files to the desired server and run them manually.

If the MySQL database is to be housed on the same physical or virtual machine as Retain, it is recommended to run the install and configuration through the install script, but it is not required or necessary. The database question in the install will not change database choices in the Retain Server Configuration and setup described in the Initial Configuration section.

MySQL Install Script

If the MySQL database system will be run on a separate server than Retain, the install and database setup may be run manually by copying both MySQL scripts and running mysqlinstall.sh on the desired server. The installation script asks whether to automatically download and install MySQL, or asks for the location of the downloaded files. The script will either download the files, or search the provided location for the installation files, download or detect them, and automatically install the packages.

```
This script will install MySQL. If you are not going to use MySQL then do not continue with this install. If MySQL is already installed, this will perform an upgrade.

Retain Installer needs MySQL install files. Please choose one of the following options:

1. I have MySQL install files already downloaded

2. I want Retain Installer to download them for me
```

After completing, if run manually, the script returns to the prompt. The script ends with instructions to manually configure the MySQL server with a root password and user. If manual configuration will be used from this point on, the commands and advice should be followed, however, if the mysqlsetup.sh script will be used, simply run the setup script to continue to automatically setup and create an appropriate database for Retain.

MySQL Setup Script

When the MySQL setup script runs, it first checks if MySQL is running, and if it is, shuts it down. After MySQL is verified to be shutdown, the configuration begins with a prompt for the Retain database location.

```
MySQL Configuration

Please specify the location where the Retain database will be stored.
The MySQL database used by Retain will be created in this location (eg: /data/mysql), or press Enter to use default location at [/var/opt/beginfinite/mysql]: _
```

The destination folder must be specified, as there are no defaults.

```
MySQL Configuration

Please specify the location where the Retain database will be stored.
The MySQL database used by Retain will be created in this location(eg: /data/mysql): ^[[C]

Create MySQL Database in /mysql-archive

Create MySQL Database in this location [yes]: _
```

Confirm the location specified and displayed by typing 'yes' and hitting 'enter', or input 'no' to change it. Next, the script prompts for the desired root username and password. The password must be confirmed.

```
MySQL Root User Setup

Please enter the desired MySQL root username [root]: root
Please enter the desired MySQL root password:
Please verify the password:

Setting MySQL Root User and Password. Write down and save this User and Password for personal reference!

Press Enter to continue . .

-
```

Enter and confirm the desired root username and password. This user is not the username which will be used by Retain to connect to the MySQL server, that user is setup later. This user will be the root user for MySQL. Hit 'enter' after recording the provided password. Do not lose this password - without it, your MySQL Database Server will not be usable, as Retain does require the root password when you are first installing Retain and when making a major upgrade

The MySQL database name, database user, and password specified here are all Retain specific. This database will house the Retain configuration as well as the stored messages and attachments. The database

user specified will only have rights to the Retain database, and is required by Retain to connect to and administer the database.

```
MySQL Database/Database User Setup
Please enter the Database Name [retain]:retain
Warning: Using a password on the command line interface can be insecure.
Please enter the retain Database username [retain]:retain
Please enter the desired MySQL retain password:
Please verify the password:
MySQL Database = retain
Database User = retain
Database User Password = [hidden]
Write down and save this user and password.
It will be needed for the Retain Wizard.
Create Database with these values [yes]: yes
Warning: Using a password on the command line interface can be insecure.
Warning: Using a password on the command line interface can be insecure.
Warning: Using a password on the command line interface can be insecure.
MySQL DB Creation successful
```

It is recommended to call the database 'retain', though the username and password may be anything desired.

The Retain database and database user specified and listed for confirmation. Do not lose this information, as it is required for the initial Retain Server setup. When the information required is acceptable, type 'yes' and hit 'enter' to allow the setup to complete. The mysqlsetup.sh script will finish, and return the terminal prompt if run manually, or will continue with Retain installation if started by the Retain installation.

MySQL 5.6 Manual Installation Procedure on Linux

Applies to: SLES 10/SLES 11

- 1. Log in as root
- 2. Get MySQL 5.6
- 3. Open a command-line terminal (like a DOS box in Windows)
- 4. Go to the directory where the RPM is.
- 5. For each rpm you downloaded, type rpm -ihv <rpmfilename>.rpm
- 6. Create a directory where you want your data stored. Type: mysql install db --datadir=<datadirectory>
- 7. Start MySQL:

MySQL is a standard Linux service, with a script stored in /etc/init.d. It is configured to run automatically on restart.

The Retain installer on Linux has automated the remaining steps of MySQL configuration. The following steps remain here as a courtesy for those who wish to manually create and add privileges

- to a database. No further configuration is needed if you are using the automated database creation provided in the installer.
- 8. From the command line, type: /etc/init.d/mysql <command> --datadir=<datadirectory> where available commands are stop, start, and restart, as appropriate.
- Create a password type: mysqladmin -u root password '<yourrootpassword>'
- 10. Log in to MySQL by typing: mysql --user=root --password=<yourrootpassword>

You are now ready to create the database. See below.

Create a new database for Retain

MySQL

Applies to: All platforms

- 1. Log in to MySQL, as in the last step of the installation procedures.
- 2. Enter the command:

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

- * This is critical. Not storing the database in UTF-8 format guarantees irreparable corruption for non US-ASCII characters. *
- 3. If you make a mistake, enter this command: DROP DATABASE retain; ...and repeat step 2.

Create a user for retain and grant rights.

Retain needs a user account created that has full rights to the Retain database. (It is strongly recommended that the 'root' account is not used.) While logged in to MySQL, enter the following commands. Replace (UserName) and (Password) with the desired user name and password:

```
GRANT ALL PRIVILEGES ON retain.* TO '(UserName)'@'%' IDENTIFIED BY '(Password)';

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

Example:

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

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

(Both commands are needed. The '%' is a wildcard. If you want to restrict the connection to a specific ip address, put that address instead of the '%' sign.)

Further tuning

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

^{*} MySQL is not set to run on restart with SLES 10, and 11. Change this with the run level Editor in YAST.

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

The following appears in the Retain Server log:

2008-10-27 00:00:59,786 [TP-Processor2] ERROR com.maintainet.dao.HibernateUtil - SCHEMA UPDATE FAILURE: Something went wrong during Schema Update - contact tech support immediately 2008-10-27 00:00:59,786 [TP-Processor2] ERROR com.GWAVA.utils.ErrorHandle - reportError: SchemaUpdate :: EXCEPTION : org.hibernate.exception.GenericJDBC-Exception: could not get table metadata: Audit org.hibernate.exception.GenericJDBCException: could not get table metadata: Audit

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

mkdir /var/lib/mysql/tmp

chown mysql:mysql /var/lib/mysql/tmp

add the following line into the [mysqld] section of /etc/my.cnf: tmpdir = /var/lib/mysql/tmp Restart the Server

SQL Server 2008 R2, 2012, 2014

- 1. Install as default instance, with Latin 1_ General_C1_AS as standard encoding.
- 2. Enable SQL Server Authentication in addition to Windows Authentication
- 3. Once installed, Verify server is listening on port 1433 (telnet ip address 1433)
- 4. Ensure TCP/IP is enabled Run SQL Server Configuration Manager
- 5. Expand SQL Server Network Configuration
- 6. Protocols for MSSQLSERVER
- 7. Enable TCP/IP
- 8. Run SQL Server Management Studio, connect to Server
- 9. Under Security/Logins, right click and select "create Login, (for example Retainuser) that uses SQL Server authentication. Assign a password.
- 10. Right click on Databases, and create a new database named Retain and assign Retainuser as the owner. (There are other ways to grant the permissions but this is easiest)

ORACLE 10/11g

Retain supports the usage of Oracle 10 and 11 on all platforms. However, it is the customer's responsibility to employ an experienced Oracle DBA to install, maintain, and tune Oracle.

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

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

Using Oracle with Retain consists of:

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

Installation

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

TCPIP Listener

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

Database Creation

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

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

connect "SYS"/"&&sysPassword" as SYSDBA

set echo on

spool C:\oracle\product\10.2.0\admin\Retain\scripts\CreateDB.log

startup nomount pfile="C:\oracle\product\10.2.0\admin\Retain\scripts\init.ora";

CREATE DATABASE "Retain"

MAXINSTANCES 8

MAXLOGHISTORY 1

MAXLOGFILES 16

MAXLOGMEMBERS 3

MAXDATAFILES 100

DATAFILE 'C:\oracle\product\10.2.0\oradata\Retain\system01.dbf' SIZE 300M REUSE AUTOEXTEND ON NEXT 10240K MAXSIZE UNLIMITED

EXTENT MANAGEMENT LOCAL

SYSAUX DATAFILE 'C:\oracle\product\10.2.0\oradata\Retain\sysaux01.dbf' SIZE 120M REUSE AUTOEXTEND ON NEXT 10240K MAXSIZE UNLIMITED

SMALLFILE DEFAULT TEMPORARY TABLESPACE TEMP TEMPFILE 'C:\oracle\product\10.2.0\ora-

data\Retain\temp01.dbf' SIZE 20M REUSE AUTOEXTEND ON NEXT 640K MAXSIZE UNLIMITED

SMALLFILE UNDO TABLESPACE "UNDOTBS1" DATAFILE 'C:\oracle\product\10.2.0\ora-

data\Retain\undotbs01.dbf' SIZE 200M REUSE AUTOEXTEND ON NEXT 5120K MAXSIZE UNLIMITED CHARACTER SET AL32UTF8

NATIONAL CHARACTER SET AL16UTF16

LOGFILE GROUP 1 ('C:\oracle\product\10.2.0\oradata\Retain\redo01.log') SIZE 51200K,

GROUP 2 ('C:\oracle\product\10.2.0\oradata\Retain\redo02.log') SIZE 51200K,

GROUP 3 ('C:\oracle\product\10.2.0\oradata\Retain\redo03.log') SIZE 51200K

USER SYS IDENTIFIED BY "&&sysPassword" USER SYSTEM IDENTIFIED BY "&&systemPassword"; spool off

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

User Account Creation

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

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

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

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

RESOURCE CONNECT

And the following database privileges to the Retain database:

CREATE PROCEDURE

CREATE SESSION

CREATE SYNONYM

CREATE TABLE

CREATE TYPE

CREATE VIEW

Tuning

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

Postgres 9

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

- 1. Connect to PostgreSQL using its psql utility.
- 2. Create a new user for use by Retain. In the example below, we are creating a user named "Retainuser" with a password of "Retainpassword":
 - CREATE ROLE Retainuser LOGIN ENCRYPTED PASSWORD 'Retainpassword'; The psql utility should respond with "CREATE ROLE."
- 3. Create a new database for use by Retain. Assign the user created in the previous step as the owner:
 - CREATE DATABASE Retain WITH ENCODING='UTF8' OWNER=Retainuser; The psql utility should respond with "CREATE DATABASE."
- 4. If Retain Server is running on a different machine, you'll need to configure PostgreSQL to permit access:
 - By default, PostgreSQL only permits local "loopback" connections. To allow connections from other machines, add the following to the postgresql.conf configuration file: listen addresses='*'
 - Specify the user and IP address that is permitted to connect to PostgreSQL by editing the pg_hba.conf configuration file. The following example grants the user "Retainuser" access to the database "Retain" from IP address "192.168.2.2":

```
# TYPE DATABASE USER CIDR-ADDRESS METHOD
# IPv4 local connections:
host Retain Retainuser 192.168.2.2/32 md5
```

• Restart PostgreSQL to activate these changes.

Linux Installation

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

- Apache 2.2.x
- SQL server with appropriate database created (Retain can automatically install MySQL, see notes below)

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

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

The install will complete the following tasks:

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

Express Install

The express installer is a tool utilized by support, testers, and for demonstration purposes. The tool was designed to only install the basics and is insufficient for production work. The Express install automatically configures a MySQL database, installs a fully preconfigured Server, worker, and Message Router. This configuration is not sufficient for production, and the entire system should be uninstalled before installing a production system.

The express install runs with little to no input from the user. The installer creates:

- SQL Database (Follow the prompts)
- One preconfigured Retain Server
- One Retain worker
- One Message Router
- Installer also installs Tomcat and Java, tunes the Java memory for a small system and sets the home variables for Retain.

The Retain Router is not meant for production. Do not use it to install on a production server. For more information, see:

http://support.gwava.com/kb/?View=entry&EntryID=2576

- 1. Download the retain installation package
- 2. Extract the zip file
- 3. Open a command Terminal and navigate to the directory where the extracted files are
- 4. Run the "chmod +x *.sh" command to make the files executable
- 5. Run the installer. "./RetainInstallExpress.sh"
- 6. Accept the license
- 7. Tomcat and Java installation will proceed automatically
- 8. Passwords are required for MySQL, SQL database, and Retain Server setup Specify the desired passwords at the prompts. These will be used for the following users:

- a. MySQL user "root" (MySQL admin user)
- b. MySQL user "retain" (Retain database admin user)
- c. Retain Server user "admin" (Retain Server administrator)
- 9. Installer will automatically install the MySQL 5.6 for the user

Upon completing the installation, the Retain Server is already configured. Open a browser and go to http://<Retain Server IP>/RetainServer and login as 'admin' with the password provided.

Retain Install

(If MySQL is the SQL server of choice, an installation and configuration wizard is provided which installs, configures, and creates the appropriate database for Retain. All that is required is to point to a downloaded install .rpm for the MySQL server and client packages, then follow the prompts. For more information on MySQL installation, or configuration and creation of an appropriate database for Retain, see the 'SQL Database Guides section'. The SQL server does not need to be on the same server, though for small systems having the SQL server local to the Retain Server simplifies the setup. To create the database manually, follow the instructions in the 'SQL Database Guide Section'. Otherwise, download the MySQL installation files and the Retain source files to the same machine, extract the Retain source files, and browse into the Retain source directory. Enter the following:

- 1. On the SQL database of choice, create a database for Retain and assign a user
- 2. Unzip the install files to a work directory. (Your distribution directory, e.g. /root/Retain)
- 3. Open a terminal session, such as Gnome terminal or Xterm
- 4. Navigate to the distribution directory containing the unzipped install files (e.g. cd /root/Retain) chmod +x *.sh(This command flags the install scripts to be executable)
- 5. Install Retain. Enter the following command:
 - ./RetainInstall.sh

(This commands install Retain-Tomcat and initialize the Retain installation which will install the embedded Tomcat as well. Consider removing any existing Tomcat installation. The Retain Install also can configure a MySQL database server, if MySQL is the database of choice for the Retain Server. The install will NOT configure MySQL for anything beyond defaults for working with Retain Server and full administration, memory tuning, and configuration for anything beyond default setup and database creation is still required.)

- 6. Read and accept the license agreement
- 7. Select upgrade options. There are two options: upgrading from 1.x or 2.x, or a new install / upgrade from 3.x
- 8. Select the features which will be installed on the machine. (Retain Server, Stubbing Server, Reporting and Monitoring Server, Worker, and Message Router.) Once the Message Router has been specified, the install will begin
- 9. If upgrading, the installation will detect whether Retain currently exists, and prompt for upgrade options. NOTE: If components already exist on the system, but new features are being added, all new and existing features must be selected for install, AND the upgrade must be selected with the first option: "Upgrade Retain, Preserving configuration" must be selected.
- 10. Installation completes and returns the prompt. Check the status of Apache 2 and Retain Tomcat. (They should both be 'running'.)
- 11. Java is checked, and automatically installed if necessary.
- 12. The SQL database is selected. if MySQL is to be used, the install script may automatically install and create the MySQL database. Follow the prompts.

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

service apache2 status
service retain-tomcat7 status

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

Custom Storage Path Rights

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

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

In a console, type:

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

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

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

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

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

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

Now check the directories for ownership.

l /retain

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

Tomcat now has rights to the storage location.

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

Windows Installation

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

- Apache 2.2.x or IIS
- SQL server with appropriate database created (Retain can automatically install MySQL, see notes below)

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

Apache or IIS

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

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

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

The install will complete the following tasks:

- If running/installing MySQL on the Retain Server, MySQL will be configured and a database created
- Java Development Kit 1.8 (JDK) will be checked/installed.
- Retain Tomcat 7 and any existing Retain installation will be stopped
- Retain will be installed
- Retain Tomcat 7 is restarted and set to run as a system service

Express Install

The express installer is a tool utilized by support, testers, and for demonstration purposes only. The tool was designed to only install the basics and is insufficient for production work. The Express install automatically configures a MySQL database, installs a fully preconfigured Server, Worker, and Message Router. This configuration is not sufficient for production, and the entire system should be uninstalled before installing a production system.

The express install runs with little input from the user. On Windows, the express install requires IIS to be installed and running. Everything else is installed automatically. The express installer installs:

- Tomcat and Java, tunes the Java memory for a small system and sets the home variables for Retain.
- MySQL and creates the Retain database (Follow the prompts)
- One preconfigured Retain Server
- One Retain Worker
- One Message Router

The Retain Router is not meant for production. Do not use it to install on a production server. For more information, see:

http://support.gwava.com/kb/?View=entry&EntryID=2576

The Express install procedure is simple - follow the prompts:

- 1. Run the installer: RetainInstallExpress.exe
- 2. Accept the agreement
- 3. Define install location

- 4. Define storage location
- 5. Install Java
- 6. Install Tomcat
- 7. Install Retain
- 8. Install MySQL
- 9. Passwords are required for MySQL, SQL database, and Retain Server setup Specify the desired passwords at the prompts. These will be used for the following users:
 - a. MySQL user "root" (MySQL admin user)
 - b. MySQL user "retain" (Retain database admin user)
 - c. Retain Server user "admin" (Retain Server administrator.)

Upon completing the installation, the Retain Server is already configured. Open a browser and go to http://<Retain Server IP>/RetainServer and login as 'admin' with the password provided.

Retain Install

NOTE: Be sure the server's clock is accurate before you begin! Using NTP is recommended to keep the server's clock accurate. Once Apache and a SQL server are installed, the system is prepared to install Retain's core components. Retain is installed as follows:

- 1 On the SQL database server of your choice create the retain database.
- 2. Verify that you have Apache 2.2.x or IIS installed and working.
- 3. Unzip the downloaded Retain install package to a work directory. Almost any name will do, as long as you remember it.
- 4. Run "RetainInstall.exe". NOTE: If it already exists, Tomcat will be automatically shut down by the installer.
- 5. The installer attempts to detect any previous Retain installations. If none are found, it will ask if the system is being upgraded from a previous install. Click 'no' for new installs, or 'yes' for upgrades. For upgrades, the installation directory will need to be specified in a later step.
- 6. System requirements are displayed. Click 'Next' to continue.
- 7. Agree to the license agreement.
- 8. Select the install location.
- 9. Select the program features to install. (Retain Server, Stubbing Server, Retain Worker, Reporting and Monitoring Server, and the Message Router.)
- 10. Installer attempts to detect previous Tomcat installations. If none are found, it asks whether Tomcat has been installed. Indicate 'yes' or 'no'. Please note that the environment variable 'RETAIN_CATALINA_HOME' is set after Tomcat installation.
- 11. Pre-requirements status is displayed. click 'Next' to continue.
- 12. Java and Tomcat are installed or upgraded. (Note: If there are issues with Apache and Tomcat starting correctly, set the Java initial memory pool and the maximum memory pool to the same value, appropriate to the system and expected load.)
- 13. Select either Apache or IIS. (If apache, provide the installation directory.)
- 14. If a previous version of Retain is detected, upgrade options are displayed; Overwrite or Upgrade.
- 15. Confirm settings and select 'install'.
- 16. The product will install.
- 17. Tomcat will start when the install is complete.

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

Tomcat is run as a System Service and shows up in the system services list as "Retain Tomcat 7". (Start | Run | services.msc)

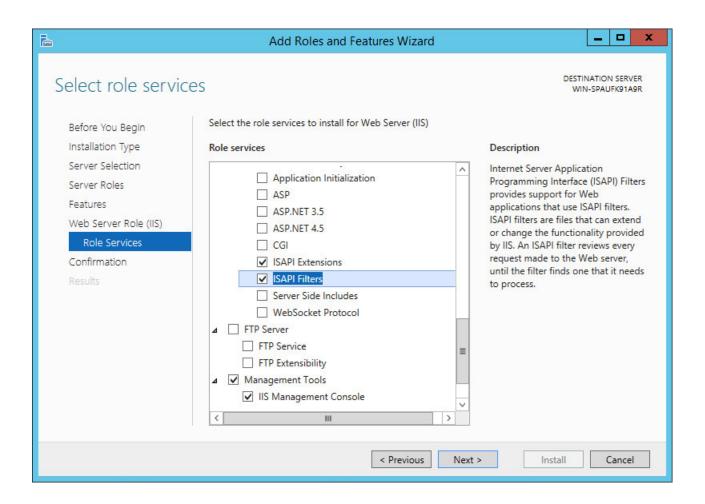
NOTE: IIS has a size limit restriction which will limit uploads, or message and attachment archiving to 30MB, by default. For more information see

http://support.gwava.com/kb/?View=entry&EntryID=2089

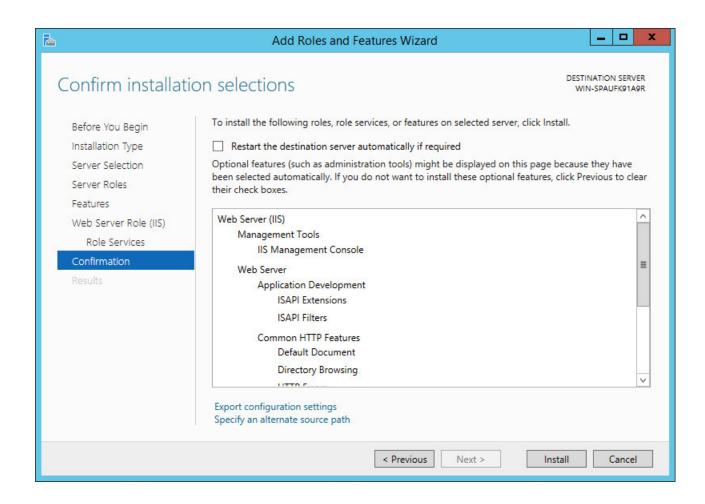
Installing ISAPI filters for Windows Server 2012 R2 and IIS 7

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

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



5. On the Select features page, click Next.



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

Post-Install Tasks

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

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

Initial Setup Overview

- Choose your storage path
- Retain prompts you for access information to your SQL Database Server
 - o You must create this database beforehand. (See install document appendix for documentation for the supported SQL servers.)
 - o Retain creates the tables, indexes, and prepares the database for use

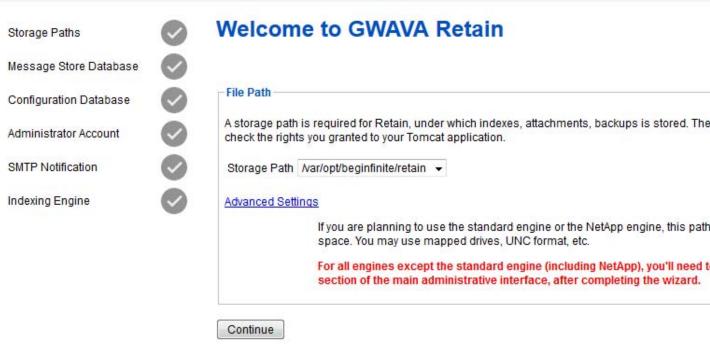
You will be prompted to provide SQL connection information for both the configuration and message store databases. For convenience and simplicity most installations store this information in the same database. The database(s) will be used to store Retain configuration, cached address books, and all saved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

- Create at least one administrator account. The first account will be the default admin account. It
 is created during the initial setup
 - o The 'admin' account is special:
 - -It has ALL admin rights
 - -It never expires
 - -It always uses "offline authentication" (authenticates using a password stored in the SQL configuration database instead of authenticating against a message system)
 - Additional accounts can be created as desired with as many rights as you wish in the management interface. User accounts may possess a mixture of administrator level and user level rights. As a result, users may have some admin level rights, but would not be considered an administrator
- Retain prompts you for the following SMTP information for notification and forwarding. (If no SMTP is present, this step may be skipped by typing the word 'skip' in the SMTP Mail Server field.)
 - o SMTP Server: Server to use for sending messages. (DNS or IP) (use DNS whenever possible)
 - o From address: Enter the address that you want to appear in the "From" field on the messages Retain sends
 - o To address: Enter the address to which Retain will send notification
 - o Username: Enter the username to use for authentication with the SMTP server for relaying
 - o Password: Enter the password for this username
- Select the indexing engine. There are two different indexing engines available for use with
 Retain: The Standard Index engine and the High Performance Index (HPI). While essentially the
 same engine, the High Performance Index is scaled and designed for high availability and rapid
 searches on very large systems. The HPI requires at least 3 separate, dedicated, index servers in
 addition to the Retain Server. The Standard Index is embedded with the Retain Server and is
 more than sufficient for small and medium systems.
 - o After initial setup, configuration of a module and creation of a profile, schedule, worker, and job are required before messages can be archived.

Initial Configuration

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





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

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

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

About Storage Paths

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

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

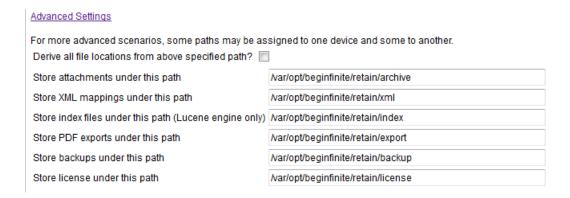
When choosing your storage system, bear in mind the following considerations:

- Storage size limits. Will storage needs ever exceed this?
- Ability to add additional storage capacity
- OS Version and kernel to support the storage system of your choice
- File system and partitioning
- Is a SAN right for you?
- Speed and reliability
- Backup/restore is your responsibility

Advanced Storage Options

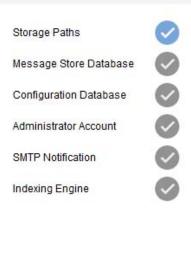
For this reason, Retain offers the advanced storage options. (Located under Server Configuration.) You can specify where you want Retain to store its information.

By default, all data is stored under one master path, specified in the Storage section. From the Retain Server main screen, you can also find this under Server Configuration under the Storage tab. If you wish to specify your storage locations in more detail, click on the "Advanced Settings".



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





Welcome to GWAVA Retain

Message Store Database Connection Parameters

To successfully create the tables the database server must be running with the retain database created, an

- Driver: Type of Database Server. MySQL 5.5 and 5.6, Oracle 10 and 11, MS SQL Server 2005/2008/20 correctly. If using MySQL, make sure you install the latest version even minor revisions contain sign
- Database Server: If the database server is not running on the same machine as Retain, change le
 DNS host name.) Including :port is legal syntax (e.g.: 192.168.2.31:3306); if not supplied, default po
- 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 us database servers require specific settings or rights to connect locally (same box setup) vs. REMOTE
- Password: Password for access.

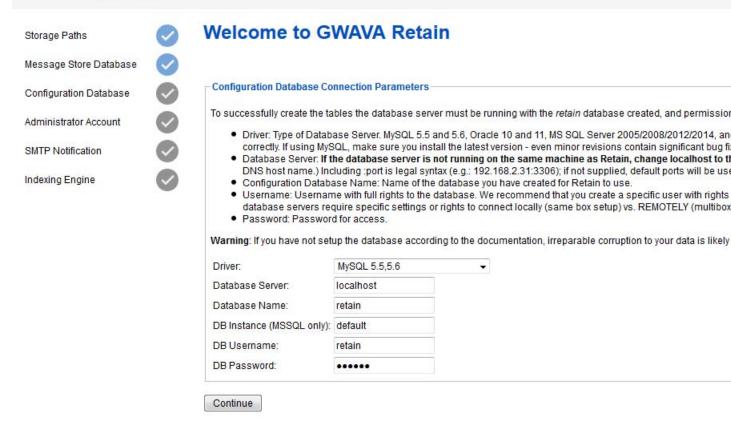
Warning: If you have not setup the database according to the documentation, irreparable corruption to your

Driver:	MySQL 5.5,5.6	▼.
Database Server:	localhost	
Database Name:	retain	
DB Instance (MSSQL only):	default	
DB Username:	root	
DB Password:		
If you are running MySQL, w DB ROOT Username:	e also need the database	server root username/password for the initial setu
DB ROOT Password:		

Continue

- 6. Next, the screen appears for entering the database connection parameters. This is the same information from the database created during the preparation phase. The connection information must be supplied for both the storage and Configuration databases, though they can be housed in the same database
- 7. The configuration database must also be specified. Unless otherwise required, the settings here should be identical to the Retain storage database, as they are usually housed together





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



Storage Paths		Welcome to GWAVA Retain				
Message Store Database						
Configuration Database		Administrator Account				
Administrator Account		This administrator account will be needed for full access to Retain. You may also create additional administrators once Keep this password secure and safe, and consider creating a backup administrator account.				
SMTP Notification		(if you do lose the account password, please contact GWAVA Technical Support to help you create a new account)				
Indexing Engine	•	Login Name: admin Password: Confirm Password:				

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



Storage Paths

Message Store Database

Configuration Database

Administrator Account

SMTP Notification

Indexing Engine

Welcome to GWAVA Retain

SMTP Notification

SMTP Notification is used for sending you progress reports, and also is used whenever a user forwards a message fro overridden on a per-job basis.

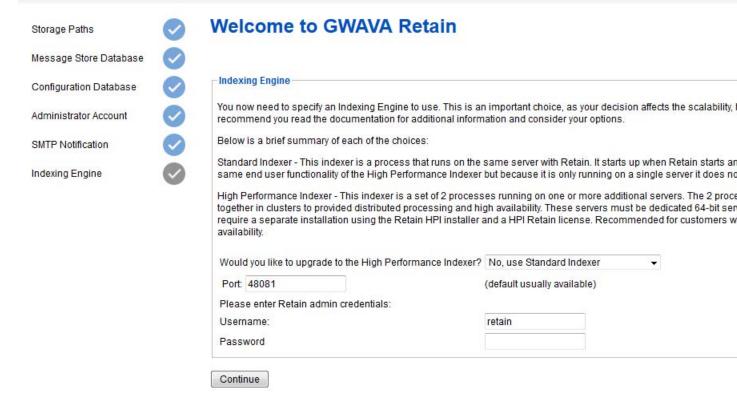
Most of these entries are self-explanatory. The TO is used for job notifications, not for forwarding. The SMTP user and pa job notifications, but you will need it if you want to permit forwarding outside your Internet domain. If connection is unend used).



Continue

10. Index Engine Selection

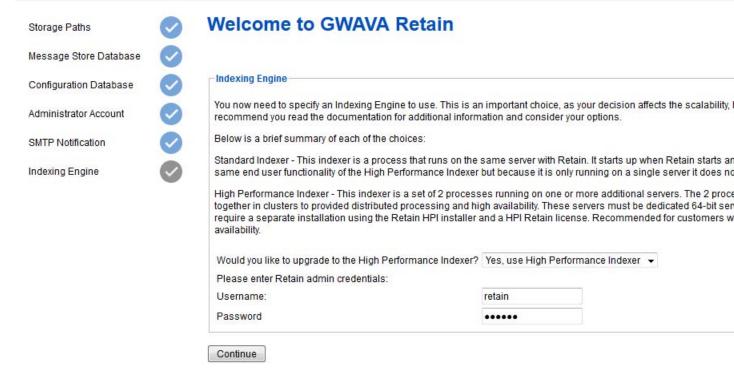




- 11. Standard is the normal, the 'standard' index engine. This engine will perform and handle most systems except the largest, is housed in the Retain Server, and requires no further configuration
- 12. High Performance Index. This index engine is designed for high load systems with very large databases. The HPI requires several separate installations and a network plan created in coordination with Support

If you wish to use HPI, it should already be installed and running.





Use the Retain administrator account name and password for the index engine connection.

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.

Once the administrator account has been created, the system will initialize and connect to the new index engine.



Storage Paths

Message Store Database

Configuration Database

Welcome to GWAVA Retain

Connecting to IndexEngine: solrcloud...

Administrator Account

Indexing Engine

SMTP Notification

 \sim



Storage Paths

vve

Welcome to GWAVA Retain

The initial setup is complete. After you login, you will need to go to Module Configuration and configure your individual con

Test Connection options in the individual modules, and then do a Address Book synchronization.

Message Store Database

 \bigcirc

accept contract and accept

Configuration Database

Administrator Account

SMTP Notification

 \checkmark

Indexing Engine

Complete

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

Front Page and First Login

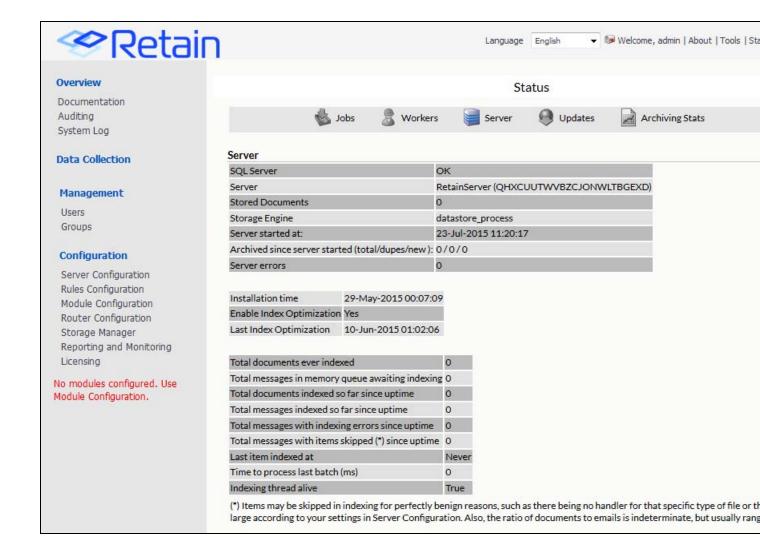
Logging In

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

You'll be brought to the login screen:



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

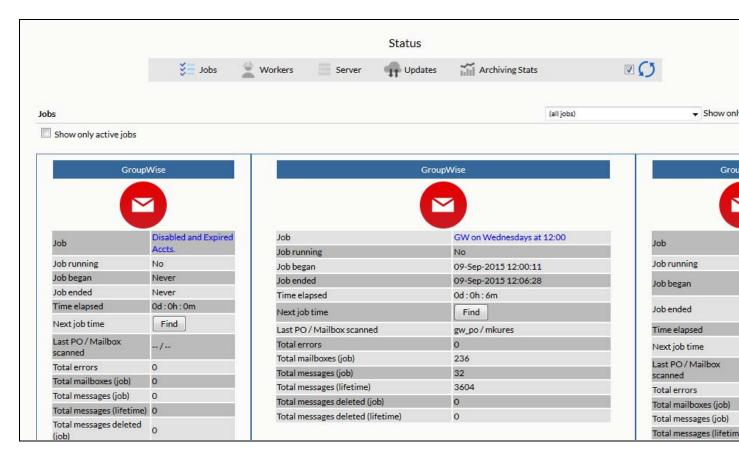


An informational screen is displayed, showing among other things the latest job activity, server health information, and software updates that are available.

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

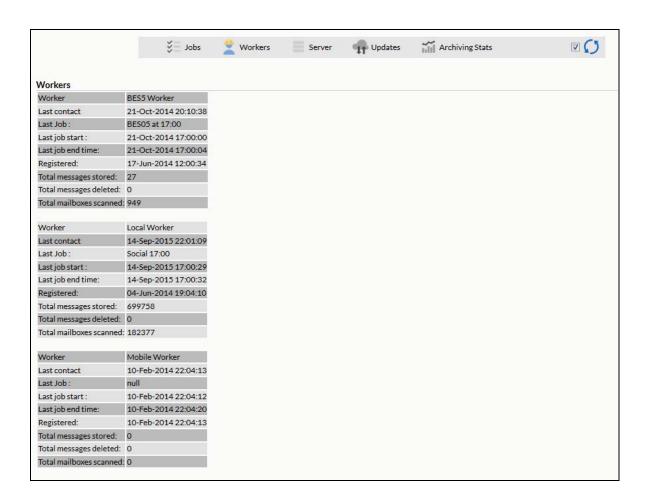
Status and Updates

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

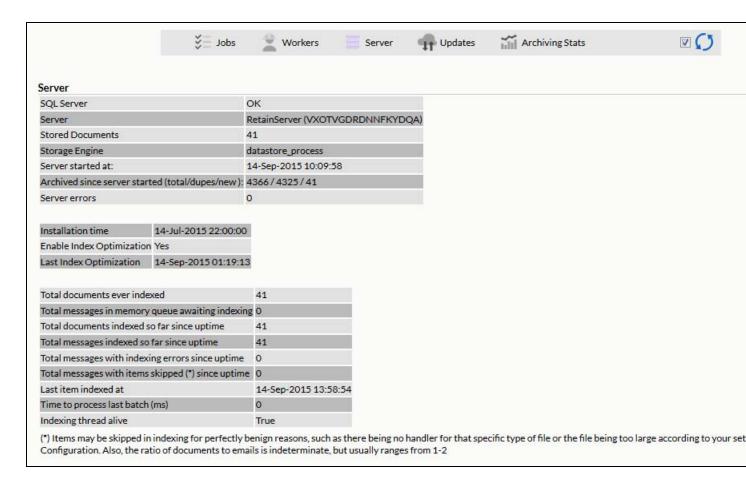


The Jobs tab, shown above, shows the jobs that are assigned to the server and their status. You also can select the "show only active jobs" check box, and only jobs that are currently running will be displayed.

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



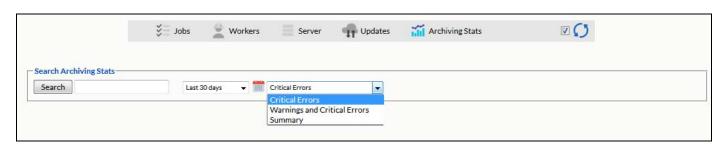
The Server tab shows the basic status of the Server with the messages archived and maintenance information. The Workers tab shows the workers connected, and the status of each.



The 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 archive stats tab is a small statistics portal which allows rapid searching of the archive statistics in the server. The search allows for a time frame, and the categories shown below: Critical Errors, Warnings and Critical Errors, and Summary. Once the 'Search' button has been pressed, a new window or tab is opened to display the search results.



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:



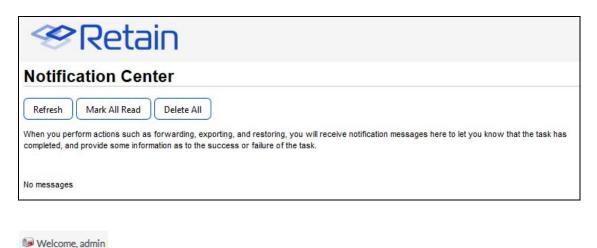
When the "Save Changes" lights up, click on the disk icon to save any changes you have made on this screen otherwise your changes will be lost.

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

Notifications

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

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



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

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

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

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

Uninstalling Retain

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

Linux

- 1. Shut Down Tomcat: "service retain-tomcat7 stop"
- 2. Shut Down Apache: "service apache2 stop"
- 3. (Optional) Remove the data storage files

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

- 4. Remove the program files
 - Can be: /opt/beginfinite/retain
- 5. Remove the RetainWorker.xml, RetainServer.xml, RetainRouter.xml, and RetainStatsServer.xml from Tomcat
 - Can be: /opt/beginfinite/retain/tomcat7/conf/Catalina/localhost
- 6. Remove the directories under tomcat/work (but NOT work itself) "rm –r /var/cache/retain-tomcat7/work/*"
- 7. Remove references to Retain in httpd.conf
 - Can be: /etc/apache2/conf.d/retain-apache.conf
- 8. Remove tomcat service file: "rm /etc/init.d/retain-tomcat7"
- 9. (Optional) Remove MySQL and Apache
- 10. If Apache and Tomcat were not removed, restart them both to verify that they load without problems.

Windows

- 1. Shut down Tomcat service
- 2. Remove Retain Tomcat 7 service

Open a terminal as admin and type: "sc delete tomcat7"

- 3. Remove the program files
 - Can be: C:\Program Files\Beginfinite\Retain
- 4. Remove the environment variables: RETAIN_CATALINA_HOME, RETAIN_HOME, RETAIN_JA-VA_HOME, RETAIN_JAVA_VERSION, and RETAIN_JRE_HOME
- 5. (Optional) Remove the data storage files
 - Can be: C:\Retain
- 6. (Optional) Remove SQL Database, SQL Server, Apache, or IIS

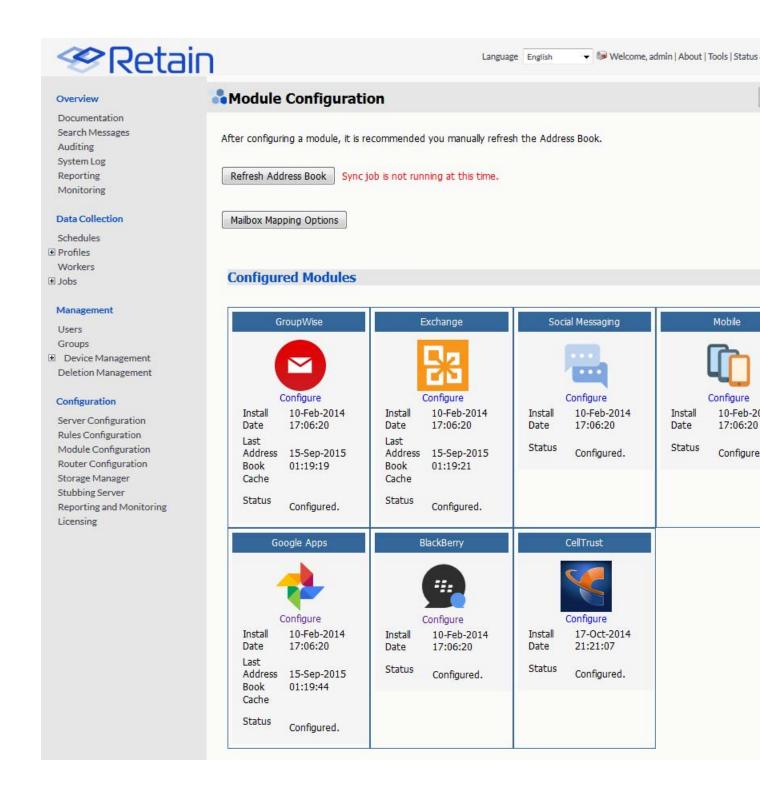
Exchange Module

Retain supports:

- A single forest Active Directory system, (Exchange and standard users)
- An Exchange Resource Forest, (One Exchange Forest linked to one or multiple User Forests)
- Office 365

Retain does NOT support multiple linked Exchange Forests. Ensure that the Exchange Settings have been configured correctly before continuing the Exchange module setup.

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

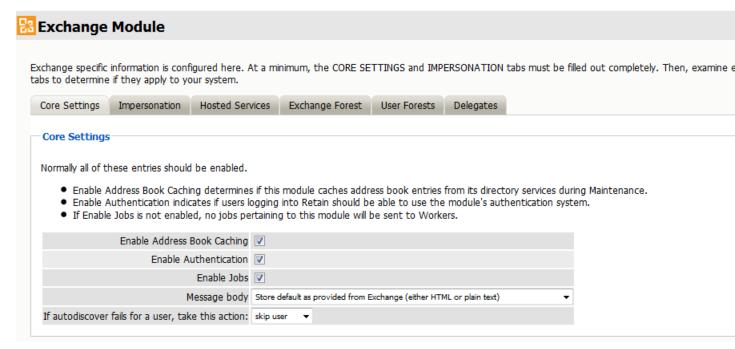


Select the 'Configure' option in the Exchange module. A new window or tab will open with the module configuration.

NOTE: Ensure that your Retain Server DNS is set to the same DNS server that your Exchange server uses. The Exchange module uses these DNS setting to auto discover critical information about Exchange stored

in Active Directory and will not function correctly unless both systems are pointed to the same DNS server.

Core Settings



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

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

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

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

The Message body allows the administrator to decide whether to store either the HTML or plain text message body, or both.

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

Impersonation

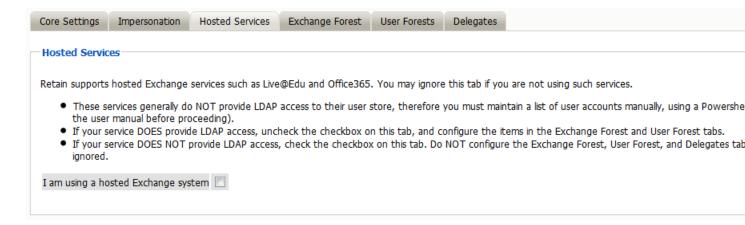
If the Impersonation and Core Settings tabs are not completely configured with the correct information, the hosted system will not be archived correctly.

Core Settings	Impersonation	Hosted Services	Exchange Forest	User Forests	Delegates	
-Impersonatio	n					
You will need to	provide this inform	mation regardless of	whether you are arc	hiving a hosted (or on premise E	Exchange system.
The impersonat	ion user MUST fulfi	Il the following requi	irements:			
 Is mail enabled. Has impersonation rights granted to all Exchange servers. 						
 Is not an Exchange Administrator (recommended) Delegation rights must be granted to any Room or Equipment accounts that you wish to archive. 						
_	_	xchange2013.ga.gv		·		
	Password ••••••			ord is set)		

Hosted Services

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

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



Mark the checkbox to enable Retain support for hosted Exchange services.

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

The Script requires PowerShell 2.0. Without this script, you cannot run jobs on Office365, or authenticate users.

Setup Process

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

Setup

- 1. Install PowerShell 2.0 (The appropriate version for the Windows system can be found here: http://www.microsoft.com/powershell) Windows 7 and 2008 R2 already come with PowerShell 2.0 (Even if the version says 1.0, really it's 2.0)
- 2. Install the Office 365 PowerShell cmdlets. Two packages are needed:
 - Microsoft Online Services Sign-In Assistant
 - Microsoft Online Services Module for Windows PowerShell
- 3. Execute them both in above order
 - Allow PowerShell script execution

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

Get-ExecutionPolicy

• The script provided by GWAVA must be run in PowerShell while running under Administrator rights.

Set-ExecutionPolicy RemoteSigned

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

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

Right click on the script and select "modify".

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

At the top you can find 3 settings: \$User, \$PlainPassword, \$ExportBasePath:

- Set \$User to the UPN of an administrator account in Office 365.
- Set \$PlainPassword to the plain text password of the administrator account.
- Set \$ExportBasePath to a directory where the two resulting csv files will be saved. If the path
 does not yet exist, it must be created manually as the directory will not be created
 automatically.

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

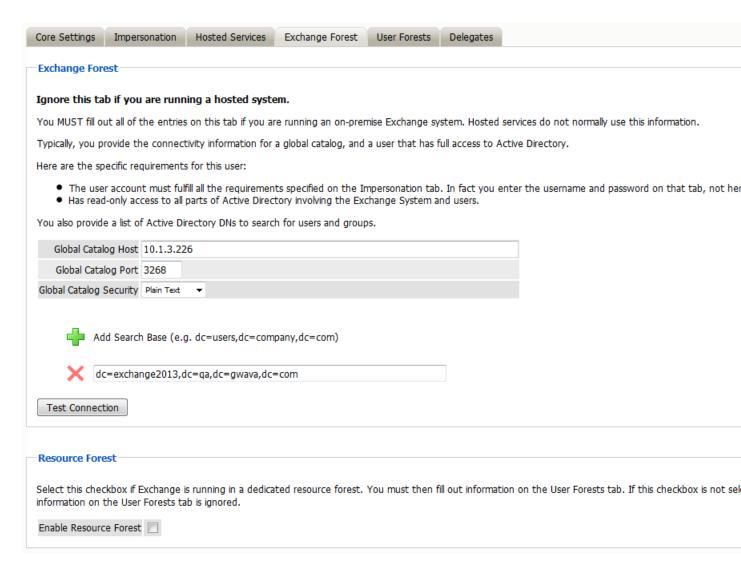
\$ExportBasePath="C:\\Temp"

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

- Execute the script by clicking the play button. This process can take a while for many users. Once the script has completed, a message indicating that the script is done should be displayed along the bottom status bar.
- 5. Automate it: to run automatically with Task Scheduler at least once a day
- 6. The frequency of automation is up to the administrator. However, the script needs to be run at least once a day, and, due to step 6, the files need to be copied to the Retain Server. To see how to utilize task scheduler to automate the script, see http://support2.gwava.com/kb/?View=entry&EntryID=2409
- 7. Get the files over to Retain Server
 - In Step 4 the ExportBasePath was set and should contain the two csv files. These files must be automatically copied to the RetainServer/WEB-INF/cfg directory once the script completes. That might require setting up a mapped network drive, or the acceptable solution for the local system. While the ExportBasePath may be mapped directly to the Retain Server, it is recommended to have the csv files copied afterwards.

Exchange Forest

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



Retain uses Active Directory extensively when integrating with Exchange. Its uses include: populating the address book, authentication, and access to the Exchange System.

There are settings required for Exchange, see the Exchange Permissions required for Retain section.

On the Exchange Forest tab, you configure all of the Active Directory information you need for an Exchange forest. There is no need to fill out any information on the User Forest tab unless the users exist in a separate forest from the Exchange Forest.

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

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

You also must provide the credentials of an Active Directory user. This user is special It must have full read rights to Active Directory, be a mailbox-enabled, user, and be granted various Impersonation and Delegation rights. More on this is discussed in the Exchange Permissions required for Retain section. The username must be in UPN format, (user principal name).

This search base, in LDAP form, must be high enough in the tree to include ALL users, groups, and servers. Multiple search bases can be specified, though it often results in a less efficient interface. These are LDAP search bases which allow Retain to resolve all users, groups, and servers of interest in the forest.

After the Search Base has been added, test the connection to ensure information and connection works. The test performs a simple login to confirm that the user exists, the Exchange Server is reachable, and that the credentials are accepted. The test does not confirm impersonation or delegation rights necessary for the Service Account.

Test Connection

Test begins....

LDAPServer connection parameters:

host=10.1.3.226,port=3268,security=none,user=retain@exchange2013.qa.gwava.com;pageSize=900;environment=key=java.naming.factory.initial;value=com.su key=java.naming.referral;value=follow

key=java.naming.security.principal;value=retain@exchange2013.qa.gwava.com

key=com.sun.jndi.ldap.connect.timeout;value=10000

key=java.naming.ldap.attributes.binary;value=msExchMailboxGuid objectguid objectsid msExchMasterAccountSID

key=com.sun.jndi.ldap.read.timeout;value=60000

key=java.naming.provider.url;value=ldap://10.1.3.226:3268

key=java.naming.security.authentication;value=simple

key=java.naming.security.credentials;value=****

Connect to and search for user:

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

SUCCESS: Connected to CN=retain,CN=Users,DC=exchange2013,DC=qa,DC=gwava,DC=com NOTE: This only tests basic LDAP authentication. If you want a mor autodiscover, impersonation, login, etc, try this link AFTER SAVING Exchange configuration:testExchange.jsp in the Util directory of RetainServer

If the test results in an error stating: "FAILURE: User doesn't exist or is not mail enabled," It indicates that the user's mailbox is unavailable. A mailbox is not required for Retain to utilize the specified user. If the user Retain utilizes does not have a mailbox, this error may be ignored. However, if the user specified does have a mailbox, this may indicate connection issues.

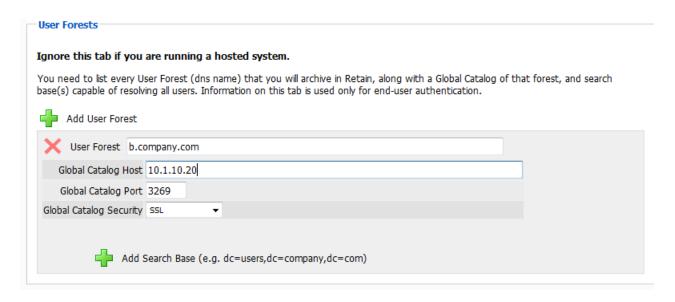
The Exchange Forest tab is the only tab required by the Server and the Worker to archive mail from the Exchange system. The User Forest tab, however, is required for Exchange systems utilizing a resource forest, to allow the end user to log into Retain.

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

If the system contains a Resource Forest, enable the checkbox on the Exchange Forest tab and save changes. If the Resource Forest checkbox is not enabled, the User Forests tab will be non-functional and

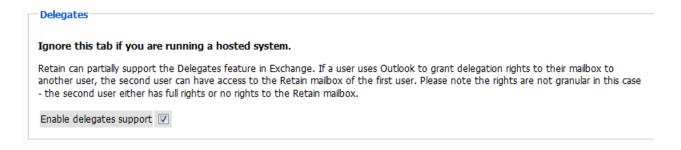
all settings contained on that tab will be ignored. The checkbox must be unchecked in a single forest Active Directory deployment, but must be checked in a multiple forest Active Directory deployment.

Check all information to ensure that it is correct and save changes, and then configure the User Forest if required.

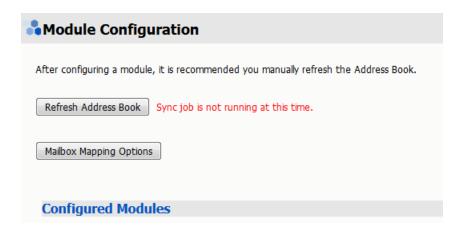


The User Forest must have an entry for each user forest attached to the system. Select the green '+' button and input the LDAP information required by the Forests' Global Catalog server: IP address or hostname, port, security, (SSL is strongly recommended), and all search bases to include all the users. No administrative credentials are required. Each end user's provided credentials will be used on login.

Save all changes before closing the Exchange Module page.



After saving changes, return to the Retain server Module Configuration page, and trigger a refresh of the Address Book.



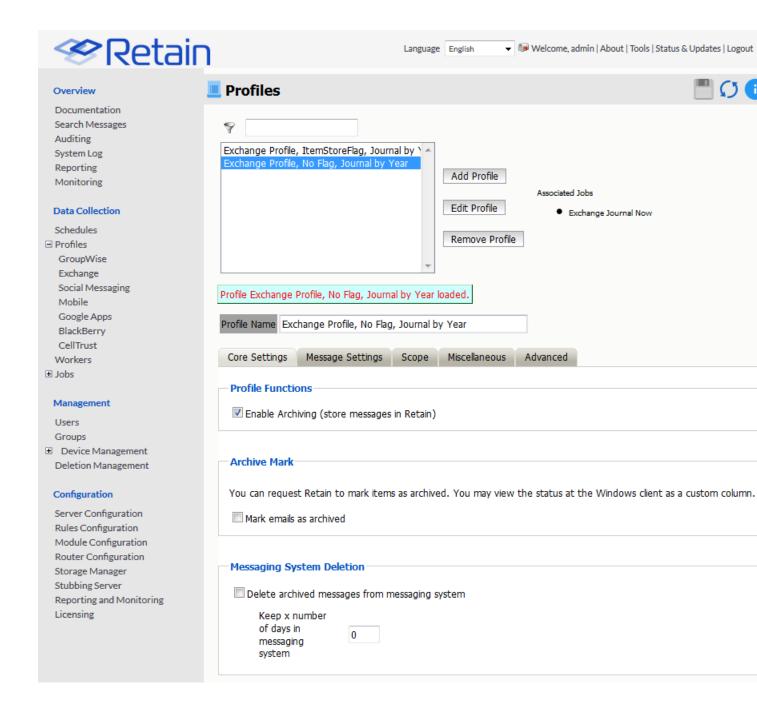
Depending on the size of the address book, it may take several minutes to return with information, but a successful configuration will return a correct address book cache date and no errors. The date should reflect the date of when the address book refresh was triggered.

The Status may show "Address Book Cache Never Run" or may list commonly misconfigured or missed items if the Refresh job fails.

Once the status is configured and the Address Book has been cached, Retain can connect to and archive messages from the Exchange server. The system is ready to have workers, schedules, profiles, and jobs configured, and those options will now appear on the main administrative interface.

Exchange Profile

After the Exchange Module has been configured, the Exchange Profile will be available for configuration. If an Exchange Profile is not configured, jobs cannot be run against the Exchange system.



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

Core Settings

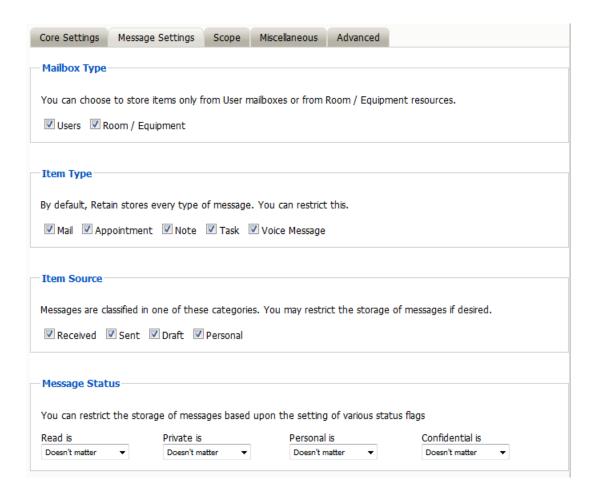
The core settings consist of an enabled/disabled option which must be enabled for any jobs based on this profile to archive anything.

Messaging System Deletion

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

Manage Settings

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



The Mailbox type specifies whether to include or exclude the available types of mailboxes. Because there can be multiple profiles and jobs, it may be advantageous to archive the Users and Room / Equipment mailboxes separately as needed and appropriate for the system.

The Item Type option specifies the different types of messages found in Exchange that can be archived, and allows the exclusion of or inclusion of the different individual types.

The Item Source option allows administrators to exclude or include messages that have not yet been sent or received, or posted.



The Message Status allows messages which have or have not been read or opened, or marked private or confidential to be archived. The different options in the drop-down menu are as shown.

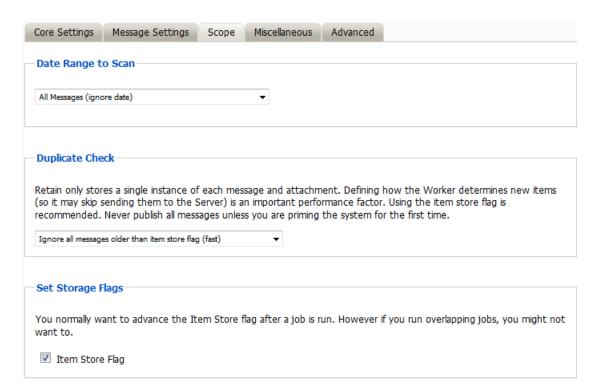
Scope

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

The Date Range instructs Retain to scan for, and archive, messages after, or before, a certain date. This is useful if only specific chunks or areas of mail are to be archived.

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

NOTE: Unlike GroupWise, Exchange does not ensure any compliance when scanning end user mailboxes; users may freely delete their email. The Item store flag does not prevent mail deletion. Only Journaling and archiving a journaling mailboxes, guarantees all items have been archived.

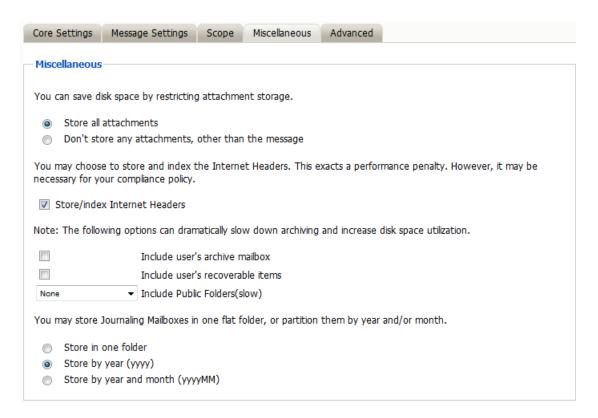


It is recommended that the first job be allowed to scan for all mail that is desired to be archived, with the duplicate check set to 'publish all messages'. Afterwards, it is recommended to have the duplicate check look for the Item Store Flag, and only check for new mail in the system.

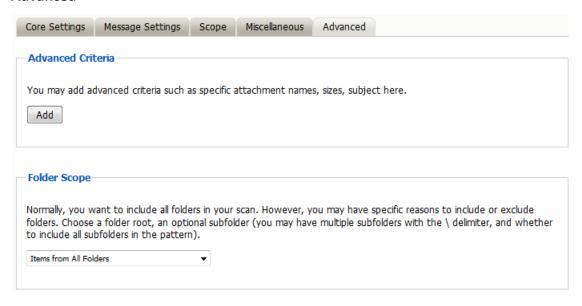
Miscellaneous

The Miscellaneous tab allows access to particular settings detailing how messages are stored and what is archived. Attachments, message information such as the internet headers, and how the data is stored and named, (by folders, year, or year and month), dictate not only the message store structure, but affect the storage size.

Miscellaneous options also allow for the archiving of the 'recoverable items'. To enable checking and archiving of the 'Recoverable Items' for compliance reasons, select the checkbox next to the option.



Advanced



If you want to be more specific as to what to dredge or not to dredge, add the criteria here. Each line will be logically AND-ed together.

Think "Dredge all items where the following is true:"

Criteria A and Criteria B and Criteria C etc

You may select based on:

Subject

Sender

Recipient

Size

Attachment Name

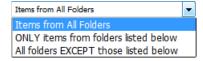
Whether they are equal to, not equal to, whether they contain or do not contain the item you specify.

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

Folder Scope

By default we dredge items from all folders. You can specify one or more inclusions or exclusions. Your choices are:

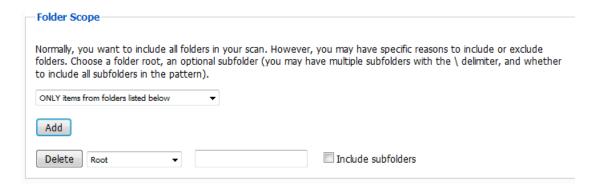
- Dredge everything
- Dredge only these listed folders
- Dredge everything EXCEPT these listed folders



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



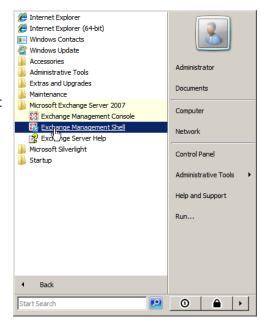
- 1. Select 'Add' to open a new selection.
- 2. Specify a System Folder (mandatory). Example: Calendar.
- 3. You may specify a subfolder of that folder (optional). Example: entering "old" would mean the folder "old" under "Calendar".
- 4. You can have multiple hierarchies under that with the / delimiter. Example "old/mail" would mean the subfolder "mail" under "old" under "Calendar".
- 5. You specify if the option includes subfolder.
 Example: If you select "old" and "includes subfolder" is unchecked, "Calendar/mail" is selected. If "includes subfolder" is checked, "Calendar/old/mail" would also be selected.



Exchange Permissions required for Retain

Exchange settings

To connect with exchange, Retain needs a user with appropriate rights. This can be accomplished by using an existing user, or by creating a new one. It is recommended to create a new user for Retain archiving. If creating a new user, ensure that the user is an active user account and that the password does not change to ensure Retain will be able to access mail without changing settings. This user is sometimes called a 'service account'. Retain calls this user the 'global catalog user'. The user created or used for Retain must be a "mailboxenabled user" with read access to see all other users, groups, resources, and Exchange Servers in the Exchange Forest. The user will be utilized by both the Retain Server and Worker for LDAP lookups in Active Directory. The Retain user also must have Exchange impersonation rights to every mailbox user on every server in the organization to be archived. The Retain user MUST NOT be a member of any Exchange Administrator group, as Exchange denies impersonation rights for all administrator accounts.



Additional permissions need to be added to the user created for Retain. The quickest way to add these rights is through the Exchange Management Shell. After creating the new user in Active Directory, open the Exchange Management Shell.

Grant Impersonation Permissions to the Retain user.

The commands required are different depending on the version of the Exchange Server. Exchange 2010, and 2013 require only one command per Exchange system to be issued, whereas Exchange 2007 requires the commands to be run on every Exchange server in the Exchange system to grant required permissions. If the Exchange system contains mixed 2007, 2010, and 2013 servers, the different commands must be completed on one server of each type.

Exchange 2010 and 2013 command:

For Exchange 2010 and 2013, the only command necessary for impersonation permissions is:

New-ManagementRoleAssignment -name ImpersonationAssignmentName -Role ApplicationImpersonation -User ServiceAccount

Where the 'Name' is a name chosen by the administrator and the 'ServiceAccount' is the name of the Retain user.

For Example:

New-ManagementRoleAssignment -Name impersonation-retain -Role ApplicationImpersonation -UserRetain

If additional Exchange servers are added to the system after running this command to grant rights to the 'retain' user, the command must be run again to grant rights to the new server.

Exchange 2007 commands:

('Retain' is the name of the Retain user, or Service Account below)

```
Get-ClientAccessServer | Add-AdPermission -UserRetain -ExtendedRights ms-Exch-EPI-Impersonation

Get-MailboxDatabase | Add-AdPermission -UserRetain -ExtendedRights ms-Exch-EPI-May-Impersonate

Get-MailboxServer | Add-ADPermission -userRetain -ExtendedRights Send-As, Receive-As, ms-Exch-Store-Admin
```

Room and Equipment Resources

To archive Room and Equipment Resources, or to restore them, the Retain user, or Service Account, must also have delegation rights. These commands must be issued manually for each Room and Equipment or resource mailbox on every relevant server. This is required for both 2010 and 2007.

These commands must be issued:

('Retain' is used here as the name of the Service Account, or Retain user, and the 'Mailbox Database' should be changed to the appropriate name.)

(NOTE: every time a new Room and Equipment or resource mailbox is added, the first command must be re-run.)

Exchange 2010 and 2013 commands:

```
Get-Mailbox -ResultSize Unlimited -Database "Mailbox Database" | Add-MailboxPermission -User "Retain" -AccessRights FullAccess

Add-ADPermission -Identity "Mailbox Database" -User "Retain" -

ExtendedRights Receive-As

Add-ADPermission -Identity "Mailbox Database" -User "Retain" -

ExtendedRights Send-As
```

Exchange 2007 commands:

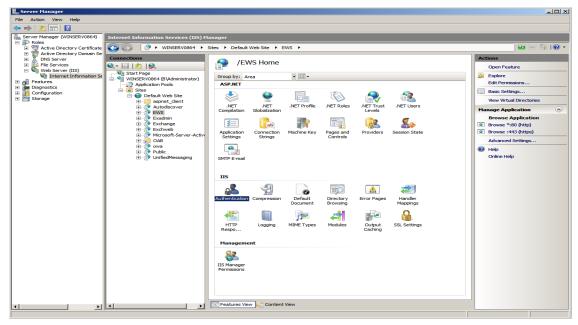
```
Get-MailboxDatabase | ForEach-Object {Get-Mailbox -ResultSize Unlimited - Database $_.DistinguishedName | Add-MailboxPermission -UserRetain - AccessRights FullAccess}

Get-MailboxDatabase | ForEach-Object {Add-ADPermission -Identity $_.DistinguishedName -UserRetain -ExtendedRights Receive-As, Send-As}
```

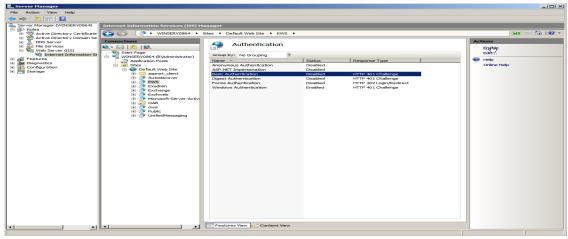
Authentication Methods

Retain requires Basic Authentication to be enabled on EACH CAS Exchange server in the system. Open "Server Manager" on Exchange server.

- 1. In left pane, expand "Roles", expand "Web Server (IIS)", select "Internet Information Services (IIS) Manager".
- 2. A new "Connections" pane opens, expand your Exchange server object, expand "Sites", expand "Default Web Site (Multiple Protocols)", select "EWS".



3. Under heading "IIS", open "Authentication" icon



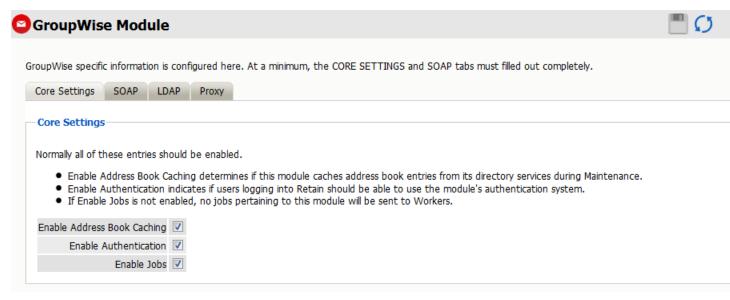
4. Select "Basic Authentication", click "Enable" in right pane. You can now close "Server Manager".

GroupWise Module

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

The GroupWise module page opens first with the Core Settings.

Core Settings



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

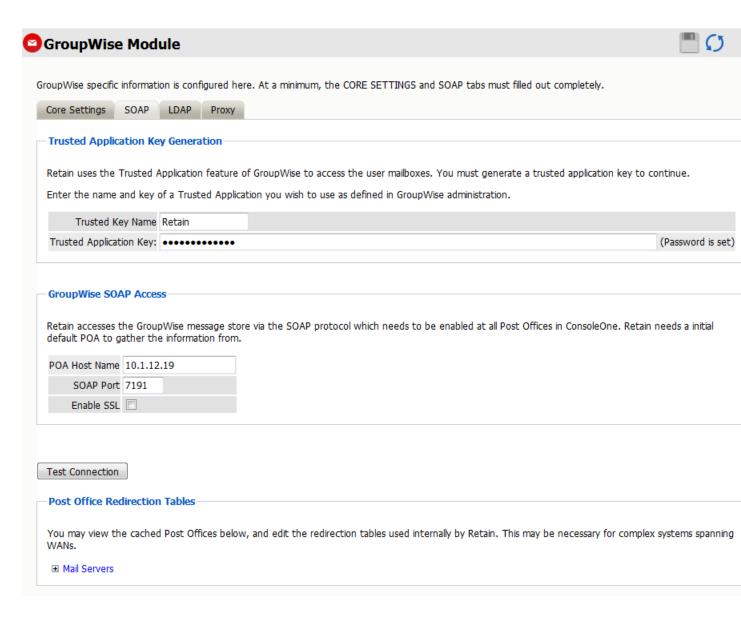
The Enable Address Book Caching function allows Retain to regularly cache the online email systems address book and synchronize it with Retain. This is critical for administration, authentication, and archiving purposes. It is recommended to cache the Address Book once every 24 hours to keep the Retain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

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

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

SOAP

Provide the POA Host Name and SOAP port. IP addresses are acceptable, but host names are preferred as IP addresses may change. SSL is supported, but comes at the price of speed, and will slow down the archiving process. Create a Trusted Application key for Retain, either manually from Console One for GroupWise 8.x and higher, or via the provided Key Generator provided and linked from the page. For instructions, see the trusted app key section.

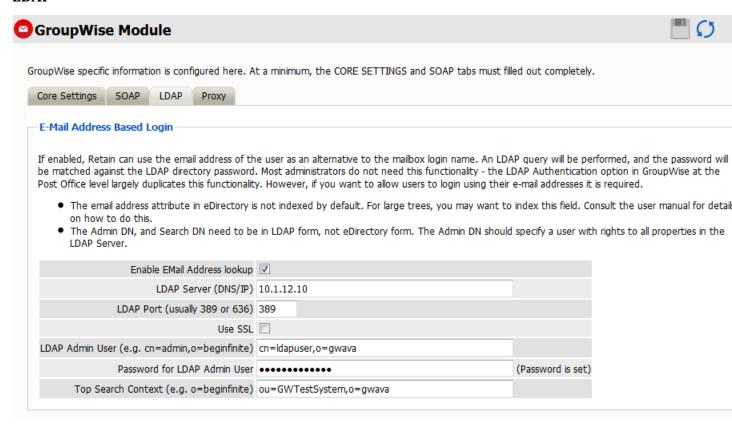


The SOAP access information must be provided, and the connection tested and verified before the system can connect. After providing the required information, click the 'Test Connection' button. The results are displayed. A successful result must be reached before Retain can archive messages from GroupWise.

Test Connection Initiate Soap Connection. Use the following parameters: host:port= 10.1.12.19:7191 trustname=Retain, trustkey=58E6.... ssl=false SUCCESS: Connected to AmandeepK For a more stringent test, be sure to refresh the Address Book Cache and then check the logs.

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

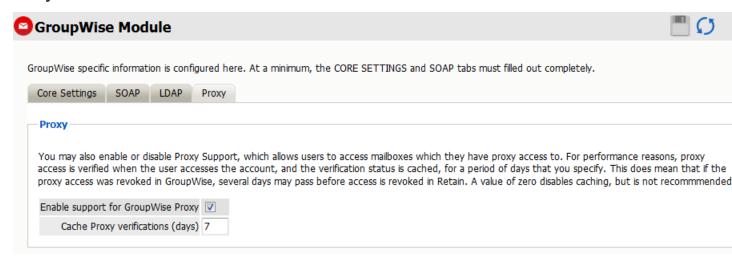
LDAP



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

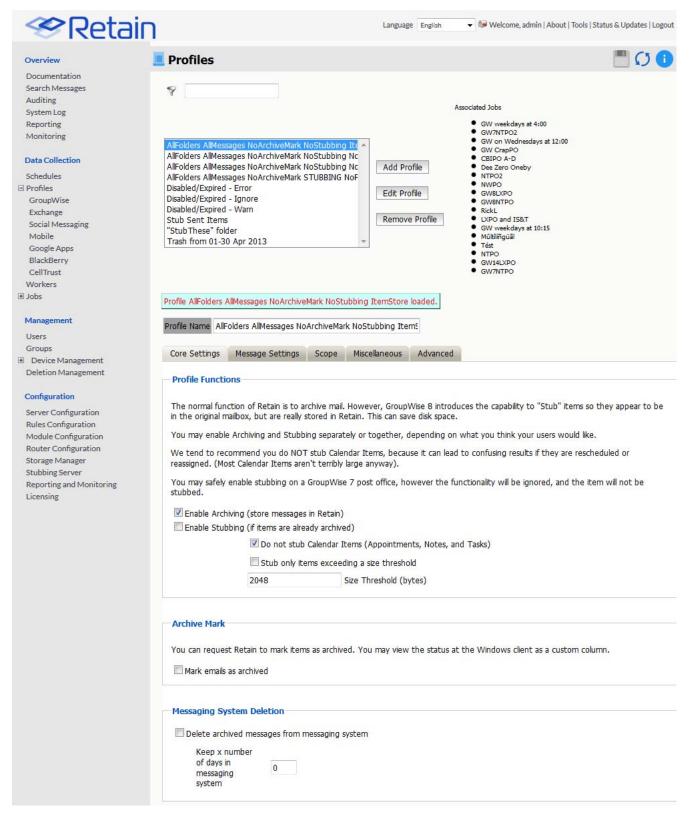
Utilizing LDAP allows users to log into their respective archives using the user's full email address. This authentication requires that the email attribute be marked indexed in ConsoleOne.

Proxy



Users who have been given proxy to another account in GroupWise may be granted access to proxy accounts in the Retain archive as well. For performance reasons, Retain caches proxy verifications for a period of days. (Default is '7' days.) Revocation of proxy access might not be reflected immediately in Retain. The caching period may be reduced or even disabled, (a value of '0' disables caching), but this is not recommended.

GroupWise Profile



Core Settings

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

Messaging System Deletion

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

Stubbing

Before enabling stubbing, make sure you understand the way stubbing works and how that will work in your system, because, once enabled, the stubbing process cannot be reversed. Stubbing, (requires GroupWise 8), allows Retain to replace archived messages and items in the GroupWise Post office with 'stubs' that link back to the full message or item in the Retain Archive. End users should not notice a difference as they have transparent access to the stubbed item through the GroupWise Post office. The stub behaves much like the native item is still in the GroupWise system.

Stubbing cannot be reversed.

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

Why stubbing is good:

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

Why stubbing may not work for your system:

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

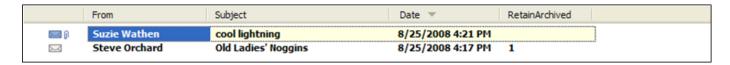
Determine the correct option for your system. If you wish to enable Stubbing, you need to install and configure the Stubbing Server as well. Refer to the install guide for your specific OS and the Stubbing Server sections. Stubbing can only be active for items which have been archived by Retain.

Calendar Items should not be stubbed unless necessary, as this can cause erratic behavior if archived items are rescheduled. If you wish to only stub large items in your Post Office, you may set a size 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.

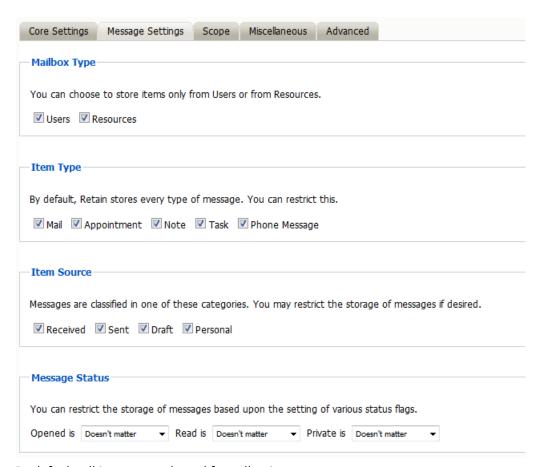


To add the column, in the client: highlight desired mailbox or folder, right-click on columns header, select More Columns, select RetainArchived, and select Add.

Message Settings:

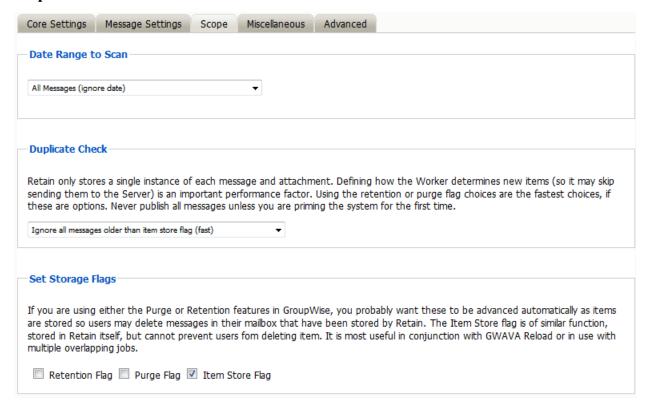
- Mailbox type: You can restrict the type of mailbox users, resources, or both
- Message Item: Select what type of items within these mailboxes to dredge
- Message Source: Select whether incoming, outgoing, draft, or posted items get dredged
- Message Status: You can decide whether to dredge an item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

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



By default, all items are selected for collection.

Scope



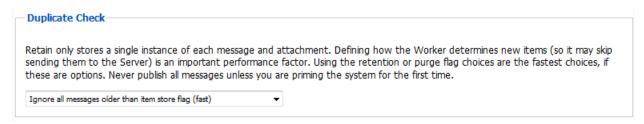
Date Range

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

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

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

Try to publish all messages (SLOW)
Ignore all messages older than retention flag (fastest)
Ignore all messages older than purge flag (fastest)
Ignore all messages older than item store flag (fast)

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 the 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 Flagfunctions very similarly to the SmartPurge flag.

• See below for a comparison

Use Item Store FlagUse an internally stored time stamp.

- Similar in function to GroupWise Retention/Purge flags.
- Determines last message stored.
- Fast option but not as fast as Retention/Purge flags.
- Ignores messages older than flag's time stamp.

Set Flags

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

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

About GroupWise's SmartPurge 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 "time stamp" 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 time stamp, stored in the user's mailbox.
- Items older than the time stamp can be deleted freely.
- Items newer than the time stamp cannot be deleted.
- A 3rd party application, such as backup software can retrieve and/or set this time stamp value. GroupWise never changes the time stamp value.
- The initial purpose of this was for backup software. Backup software can either manually or automatically, (via the SMS TSAGW component of GroupWise), advance this time stamp once all items are backed up. The backup software can also use this time stamp to "know" that all items older than this time stamp have already been backed up, and can be skipped.
- Retain can use this flag to skip archived items, and automatically advance it as well. However, in
 practice, you'll rarely want to do this, because backup software, including 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 <u>cannot</u> 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:

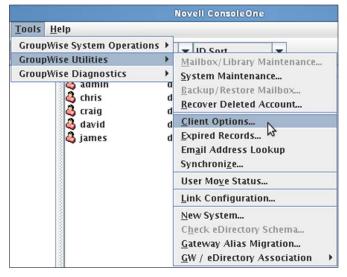
Users will never again be able to delete e-mail from the live GroupWise system.

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

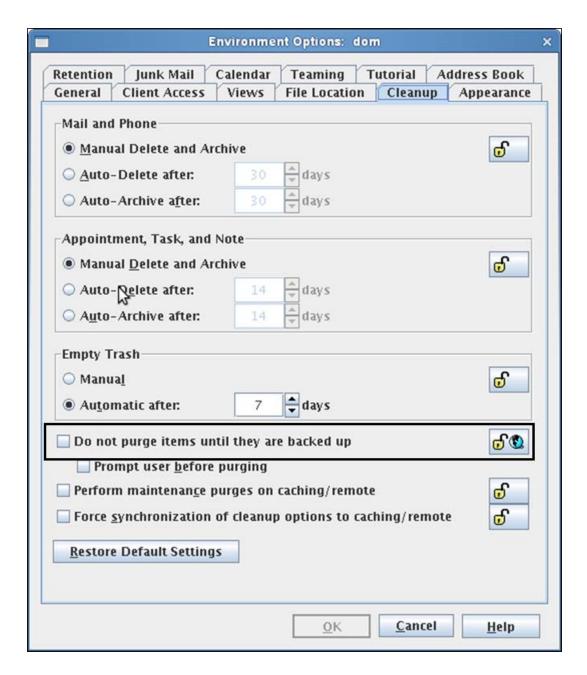
(SmartPurge Setup under GroupWise 8 is similar.) Setup for both flags is accomplished in ConsoleOne:





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

Once there, SmartPurge is set up like this:



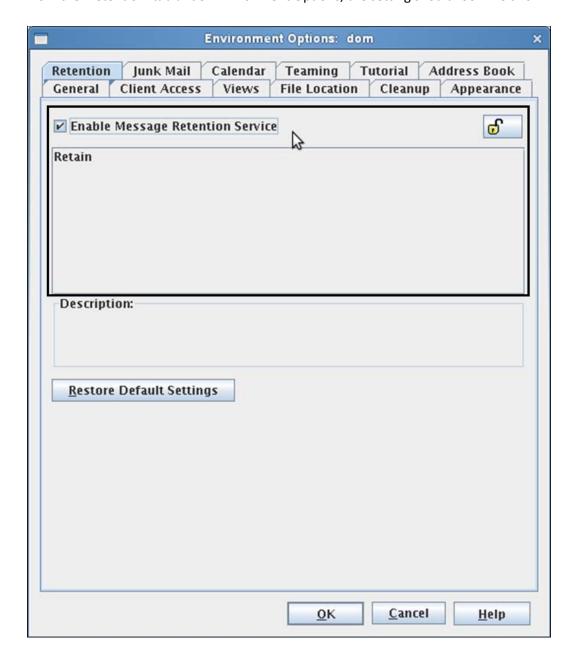
In the Cleanup tab, UNCHECK the box on "Allow purge of items not backed up" to ENABLE SmartPurge. Lock it to prevent users from disabling it.

(GroupWise 8 uses reverse logic with the purge option. GroupWise 8 lists to "Do not purge items until they are backed up." Reverse the setting for the desired effect.)

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

The Retention flag is set up a little differently than SmartPurge. (This is only used when Stubbing.) A Trusted App requesting Retention Services must be created. (Done at Retain setup time). The Retention flag is enabled on a per domain/po/user basis.

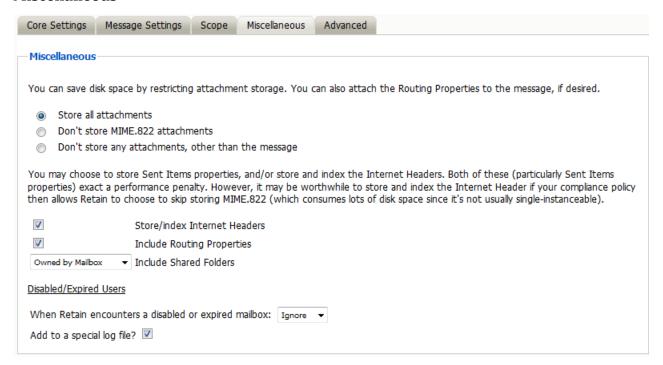
From the 'Retention' tab under Environment Options, the setting should look like this:



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

Next, check the box to enable the Retention Flag service. Lock it on.

Miscellaneous



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

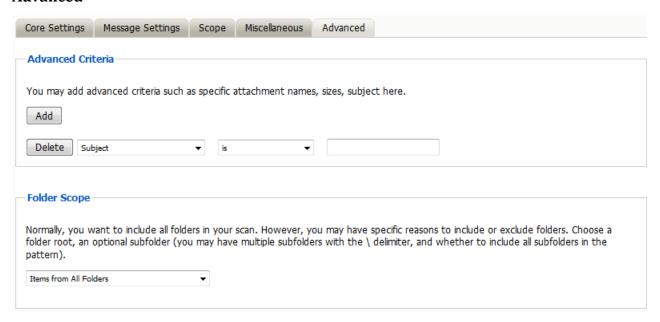
But if you don't want to store attachments, you can control this here. Note that the attachment names and sizes will still be listed when viewing the message, but the attachments will be clearly marked as unavailable.

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

Store/index Internet Headers

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

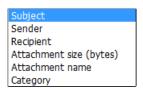
Advanced



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
- Category

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

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

Folder Scope

By default we dredge items from all folders. You can specify one or more inclusions or exclusions. Your choices are:

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

Items from All Folders

ONLY items from folders listed below All folders EXCEPT those listed below

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

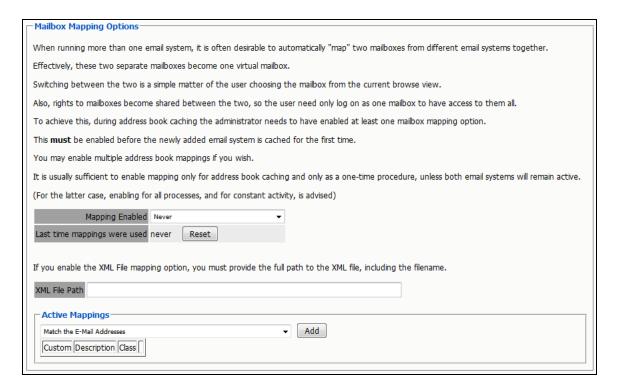
- 1. Specify a System Folder (mandatory). Example: Calendar.
- 2. You specify a subfolder of that folder (optional). Example: entering "old" would mean the folder "old" under "Calendar".
- 3. You can have multiple hierarchies under that with the / delimiter. Example "old/mail" would mean the subfolder "mail" under "old" under "Calendar".

"includes subfolder" is CHECKED, "Calendar/old/mail" would also be selected.

4. You specify if the option includes subfolder. Example: If you select "old" and "includes subfolder" is unchecked, "Calendar/mail" is selected. If

Mailbox Mapping Options

Mailbox mapping automates the merging of users' mailboxes that have both GroupWise and Exchange accounts, combining the GroupWise and Exchange archives under the same user in Retain. Retain merges the internal GUID's of accounts to reflect the same user. After mailbox mapping has been activated, an immediate caching of the address books of both systems is required. Mailbox mapping is not permanent; however, "un-merging" must be performed manually via the Mailbox Merger / Splitter tool found on the 'about' screen. If a refresh of the mailboxes is not performed before users attempt to login or other activity is present on the server, the merge may fail, leaving extra entities created in the Retain archive which must be removed manually. Furthermore, the merge option is a once-off option. Please contact support for further options.



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

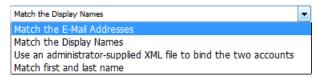


Accounts may be matched via email address, display names, or specified by an administrator supplied XML file. Most organizations maintain a naming scheme from one system to the next which results in the email address being the same in both systems. For this situation the email address match is the easiest and quickest way to map the users. For systems which have listed the first and last name of each user, but have different email addresses, the name matching is the best option. The final option for an admin-

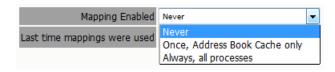
istrator provided xml file to bind the accounts together is time consuming and requires the administrator to build a list of matching addresses with UUID's in both systems.

Any mis-mapped addresses may be corrected by using the Mailbox Merger/splitter utility found on the 'About' page. see the 'About and Diagnostics Page' section for more information. A sample XML file is provided and can be download from the 'Mailbox Merger/Splitter' utility page, it is also shown below.

The Address book MUST be refreshed immediately in order to merge the users. The merge will not complete unless the Address Book is cached afterwards, and the merge cannot be re-run.



How the merge runs depends on the different users in the local system. If both message systems are to remain active, with users being added and removed from the system, it will be necessary to have mailbox mapping on all the time. Merged addresses are also required for the GroupWise to Exchange migration to function.



Mailbox Mapping Enabled

The settings here are designated for the different situations which a system migration may be in.

Never

This setting disables mailbox mapping, and is to be used in most cases.

Once, Address Book Cache only

To migrate a system, or for systems which have already been switched over, where the old system is no longer being used, but is still running.

Always, all processes

For situations where both mail systems continue to have incoming, new mail, the mapping needs to be completed repeatedly, to keep the message data mapping to be accurate for newly archived mail.

XML File

The supplied XML file is the way for an administrator to manually dictate which addresses are to be mapped or tied together in the Retain system. This process is not recommended as it is time consuming and somewhat complicated. In general, it is likely quicker to use the Mailbox Merger/Splitter tool from the 'About' page instead. Both situations require the administrator to know which email addresses are to be tied together before beginning.

The XML file requires the administrator to specify, in pairs, which addresses are to be tied. The source address is the address from which the system is being migrated.

The destination address is the address which the user is the new system, or where the data is being migrated to.

The 'removeOrphanedEntity' option tells Retain to remove any unconnected entities left over. Default for the XML file is 'true', while the default in the Mailbox Merger/Splitter tool is 'false'. By rule of thumb, if you are uncertain, don't remove it.

```
-<MergeTasks>
  -<tasks>
    -<MergeTask>
        <sourceMailbox>source0</sourceMailbox>
        <destinationMailbox>destination0</destinationMailbox>
        <re>moveOrphanedEntity>true</removeOrphanedEntity></re>
     </MergeTask>
    -<MergeTask>
        <sourceMailbox>source1</sourceMailbox>
        <destinationMailbox>destination1</destinationMailbox>
        <re>moveOrphanedEntity>true</removeOrphanedEntity></re>
     </MergeTask>
    -<MergeTask>
        <sourceMailbox>source2</sourceMailbox>
        <destinationMailbox>destination2</destinationMailbox>
        <removeOrphanedEntity>true</removeOrphanedEntity>
     </MergeTask>
    -<MergeTask>
        <sourceMailbox>source3</sourceMailbox>
        <destinationMailbox>destination3</destinationMailbox>
        <re>moveOrphanedEntity>true</removeOrphanedEntity></re>
     </MergeTask>
    -<MergeTask>
        <sourceMailbox>source4</sourceMailbox>
        <destinationMailbox>destination4</destinationMailbox>
        <re>moveOrphanedEntity>true</removeOrphanedEntity>
     </MergeTask>
   </tasks>
 </MergeTasks>
```

BlackBerry Module

The BlackBerry module allows the collection of BlackBerry phone data from the Enterprise server into the Retain system. BES 5.x, 10.x, and 12.x are supported. To archive the BlackBerry information, the BlackBerry Enterprise Server must be configured to log phone calls, PIN, BBM and SMS data. Archiving BBM Protected data is also supported through the REST API. (Archiving BBM Protected will require two additional licenses: REST API license and the BBM Protected License. See the rest API section.)

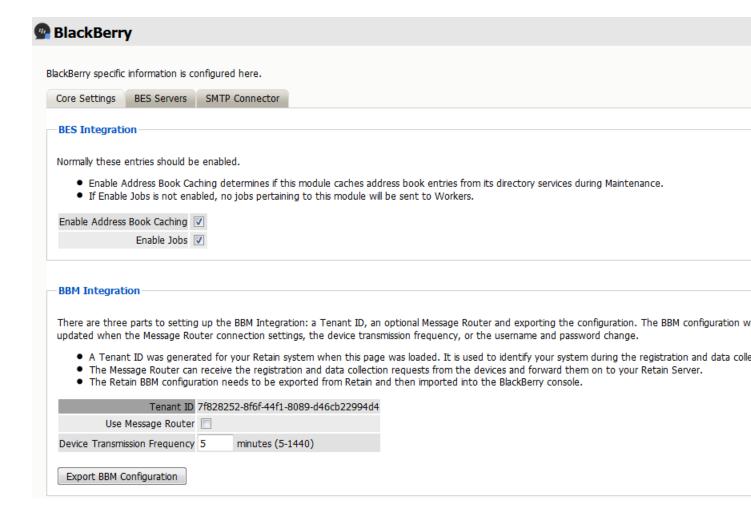
The server log files are what Retain uses to collect the information about phone activity, and as such no software or modification needs to be made to any system phones. If the BlackBerry server has not been configured to log the data, Retain will be unable to archive the information. For BES 5.x, instructions to modify the logging in the BlackBerry Enterprise Server can be found in the BlackBerry Enterprise Server help file. (Found under Start | Programs | BlackBerry Enterprise Server | Help | BlackBerry Manager Help)

For BES 10.x and 12.x, users must be set on EMM – Regulated mode and have the logging set to 'yes' for all types desired. Balanced mode users will not have logs created for them in the BES system and as a result will not be archived.

Installing the BlackBerry Module is incredibly simple. Essentially, the module needs to be enabled and configured with the BES log file location. Afterwards, the data is treated the same as any other data entering the Retain archive. Data is archived in the database under the user/phone it originated from and can be searched and viewed through the Search Messages interface.



To enable the BlackBerry Module, select the 'BlackBerry Module' from the module page and click 'configure'.

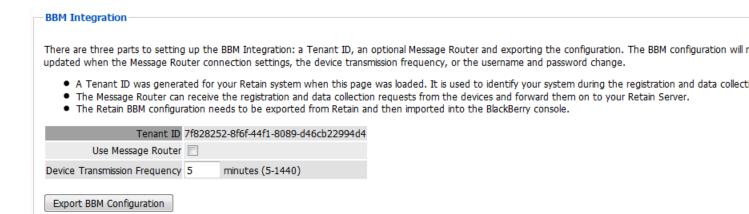


The BlackBerry module needs to be enabled on this page to make it active in the Retain system.

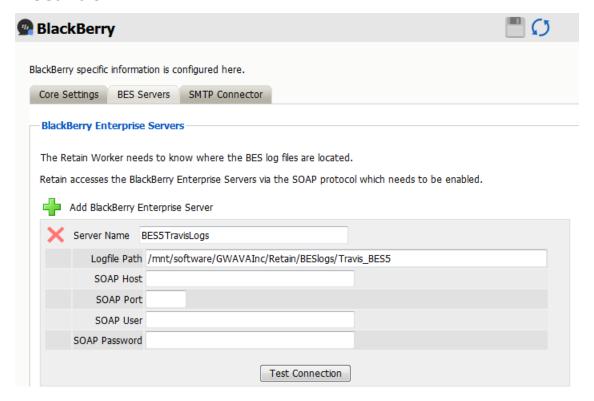
For address book caching to occur, it must be enabled on the Core Settings screen. For more information on caching, see below.

Before data can be archived, the BES Servers must be configured and a BES log file path must be provided.

If BBM Protected integration is to be used, no BES servers will be listed and no other configuration in the BlackBerry Module is required. However, the BBM Configuration file needs to be exported and uploaded to the BBM server. Please see the BBM setup and install checklist for specific information on how to configure the BBM server.



BES Servers



The log file location must be the address to the Blackberry Server log files, from the perspective of the Retain Worker. Regardless of where the worker is installed, local or remote to the Blackberry Server, it must have direct file access to the logs. This access can be accomplished through a shared folder or location, but it must be available to the Retain worker. The path must be an open and active path from the Worker to the log file location. Appropriate access through firewalls or network security must be made available if the worker is not installed on the BES system.

It is important to keep in mind that the log path needs to be the path relative to where the Retain Worker is installed. If the Retain Worker is installed on a system that is local to the log directory than a local path can be used here. If the Retain Worker is on a remote system than a UNC path must be used. If using a

UNC path to a Windows share it is necessary to ensure that the Retain windows service (tomcat) be running under a user with the proper rights and permissions to read the log data being shared.

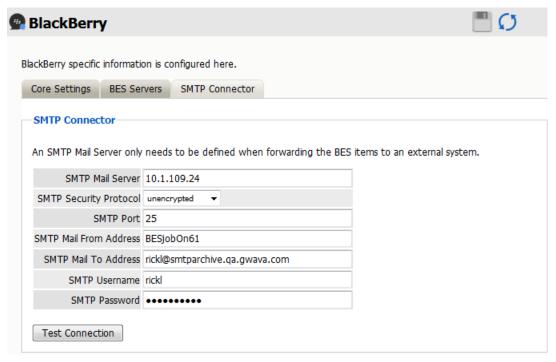
The SOAP data must also be filled-out. The SOAP connection is how the Retain Server connects to the BAS service's SOAP interface in order to cache the address list. SOAP must be enabled on every BES listed. List every BES server which the Retain system will be archiving data from. The SOAP address or host name and port which is configured must be specified for each Server. The SOAP user name and password are the credentials to the BES admin web console.

Caching

Retain will attempt to cache the address book from each server listed if address book caching is selected. Because the Blackberry Server logs, and especially the BBM Protected data does not always contain user names attached to the device ID's, Retain works to cache the address book to fill out data for messages; user name, email address, etc. Without Address Book caching, stored message data may only be identified by a device ID or phone number. Each server's configuration has a test button to help ensure that Retain can successfully connect to the server. Selecting the 'Test Connection' button will initiate a connection attempt by Retain for the server. The results will be displayed after the test has been completed.

SMTP Connector

In order to allow the forwarding of BES items out of the Retain system, the Blackberry Module needs to be configured with an SMTP connector. This connector provides an option for Blackberry data to be forwarded out of Retain.



The amount of configuration necessary for the SMTP Connector depends on the specific setup of the SMTP system. If a username and password are required for the host SMTP system, then it must be provided. If it is not required, the configuration of the username and password is optional.

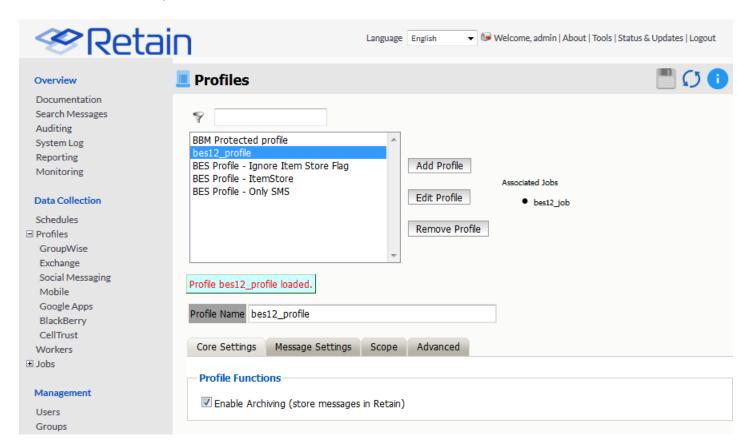
The address of the Mail Server, the SMTP Mail From Address, and the SMTP Mail To Address, must all be configured for the system to function correctly. The 'to' address is the destination address where the data will be forwarded to. The 'from' is the specified 'from' address.

In addition, the 'Enable SMTP Connector' option must be enabled in the 'jobs' configuration for the Black-berry section.

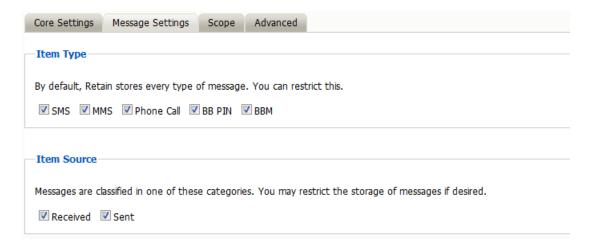
BlackBerry Profile

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

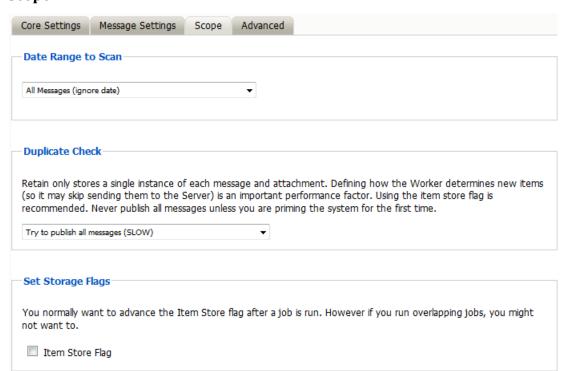
The Core Settings tab holds the enabling setting, which allows jobs connected to this profile to run. This must be checked before jobs will run.



The message settings tab can be used to exclude or include specific data types for this profile. By default all types of messages will be archived. All BlackBerry data which is selected to be archived must be logged in the BlackBerry Enterprise Server or the data will not be available for Retain to archive it.



Scope



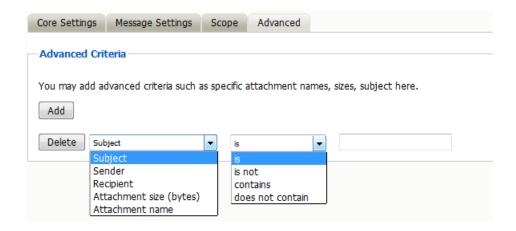
The Scope tab dictates how pervasive the attached archiving jobs are, and what kind of duplicate checking is done, or not done, in the Retain system.

The Date Range instructs Retain to scan for, and archive, messages after, or before, a certain date. This is useful if only specific chunks or areas of mail are to be archived.

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

It is recommended to have at least the first job archive everything, with the item store flag checked to archive everything possible in the logs, while setting the item store flag. Subsequent jobs should use the duplicate check and the item store flag to only archive new data afterwards.

Advanced



If you want to be more specific as to what to dredge or not to dredge, add the criteria here. Each line will be logically AND-ed together.

Think "Dredge all items where the following is true:"

Criteria A and Criteria B and Criteria C etc.

You may select based on:

- Subject
- Sender
- Recipient
- Size
- Attachment Name

Whether they are equal to, not equal to, whether they contain or do not contain the item may be specified.

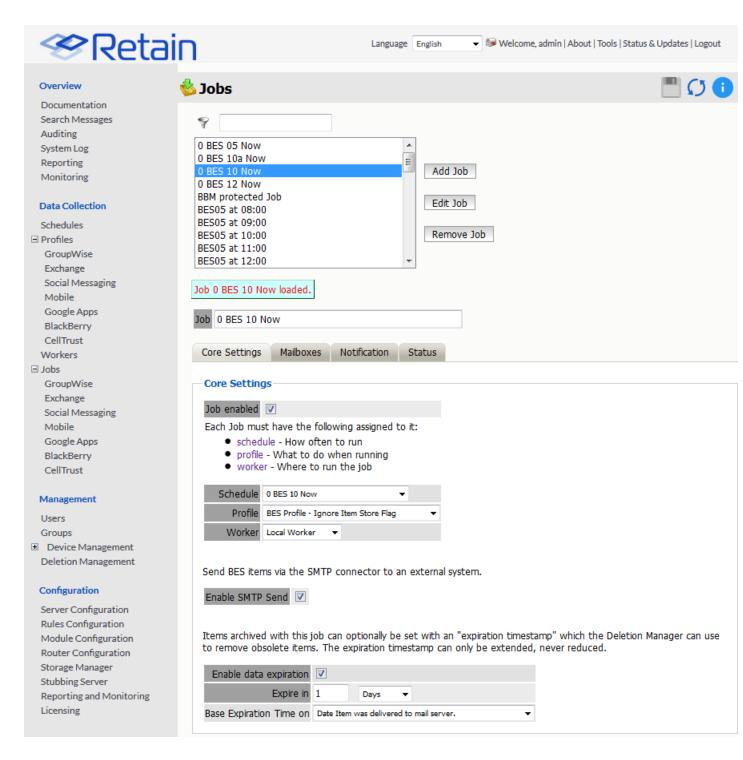
This provides great flexibility and allows you to customize dredges and retention for many different groups, or even individuals.

Blackberry Jobs

All data collection for the Blackberry module is configured through the Blackberry job interface. The jobs here combine the Profile, the Schedule, and a Worker together to archive the specified server logs. A Job must have a Profile, Schedule, Worker, and a selected server, before it can be saved or run. To create a new job, click, 'Add Job', and configure and save.

Core Settings

The core settings tab contains the basic configuration. Select the 'Job enabled' option to enable archiving. Also select a Schedule, Profile, and Worker to complete the job. If the SMTP Send is enabled, the worker will forward BES message data.



The Data Expiration is an option to place a timestamp on the data in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, Caringo, and Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one (including Retain) may delete the item before its expiration date. Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for. The base folders and criteria are specified under the custom expiration dates accessed through the 'Add' button at the bottom. All messages included in any specified folders will have a different date or be exempted from the

standard expiration date. In order to have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

Mailboxes

Core Settings	Mailboxes	Notification	Status	
Mailboxes —				
Select the BES server to be included in this job.				
BES10 via NFS	•			
⊞ Users				

The mailboxes tab contains the target users, and BES servers to be archived by the selected job.

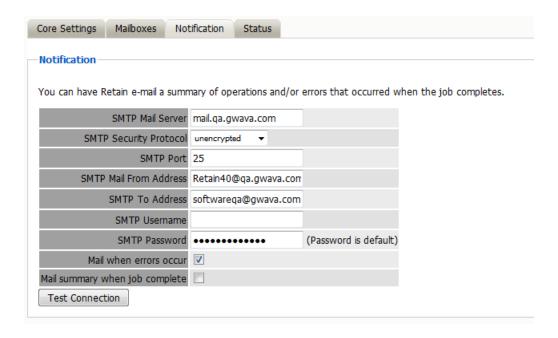
Under the Mailbox tab, select the BES server to be included in the Job. The list of available BES servers is configured from the Modules page. Once selected, all users associated or present on the BES server are included.

Once a job has been run, the user list will be populated and further configuration is possible. The job may also be restricted to specific users, specified on the 'Include Users' list, or restricted to exclude users listed on the 'Exclude' list. To add users, expand the User's option box and select the users to be archived. Users may be selected through the user search interface, spawned through selecting the 'add user' button.

If there are users listed, then the users listed in the 'Include Users' box are the ONLY included users, while those listed in the 'Exclude' window are excluded from the jobs.

Notification

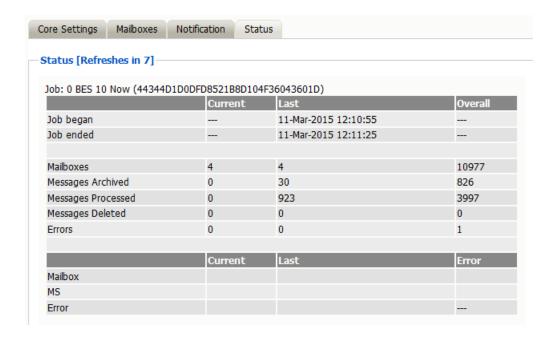
All errors, summaries, and statistics may be sent in an email notification to a desired address on job completion. If it is desired to be notified, configure the notifications under this tab. At least one of the two notification options must be selected in order for an email notification to be created.



Status

On initial creation, there will be nothing of interest shown under the 'Status' tab, as the job is under creation or has not yet run.

However, after a job has been run, or has begun, the 'Status' tab becomes a monitoring tool to report on exactly how much mail and mailboxes have been archived, along with any errors encountered.



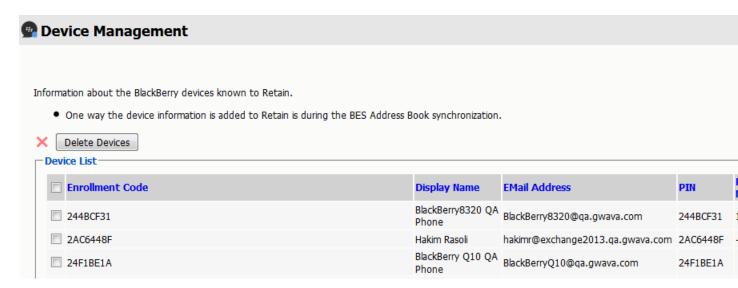
Device Management

Devices are deleted from Retain through the Device Management interface. Device Management is split into two sections: Mobile and Blackberry. In both sections, multiple devices may be selected and acted

on. In addition, all selected devices remain selected until either an action is taken or the administrator has navigated away from the page.

BlackBerry

Retain supports archiving data from both BlackBerry devices and devices running the BBM Protected system. Devices running the BBM Protected system can have their communication data archived through the BBM app, obtained from BlackBerry.



The BlackBerry device page displays all devices which have been registered into the Retain system. Native BlackBerry and BBM Protected devices are displayed together. Retain registers BlackBerry devices through log files from the BES server or through data sent by the device running the BBM Protected application. If a device is not yet registered, it is either not in the BES address book or no data has been archived for the device. The device list may be filtered for any specified content desired. In addition, devices may be selected for removal.

BlackBerry Auditing and Archiving Service for BBM

The setup for Retain BlackBerry Auditing and Archiving Service for BBM Protected (BAAS) is quite simple for Retain, as the BBM Protected application directly inputs data to the Retain BAAS system. The BBM Protected app can forward data to either the Retain Router or the Retain Server, but it will only archive to one or the other. Because the BBM Protected app communicates directly to the Retain BAAS system, Retain needs to have an open connection to the Internet. Due to this requirement, the desired connection should be placed in the network's DMZ. To protect the Retain Server, the Retain Router has been designed to be placed in the DMZ. When the Router cannot be installed in the system, the Retain Server can handle all communication, but will need to be installed in the DMZ. Whichever will be used, Router or Server, the configuration and a security certificate must be obtained before setup can be completed.

Item and Requirements Checklist:

- Enterprise ID (From Blackberry)
- Base API license (In licensing tab)
- BlackBerry license (In Licensing tab)
- REST API base License (In REST API tab)
- BBM REST API license (In REST API tab)
- SSL security certificate obtained from a trusted third party certificate authority for either Server or Router (whichever is to be used).

Pre-setup tasks:

- Network structure plan (Where to place the Retain Server, and or to use the Router.)
- Retain Server setup and configured or Server and router configured
- Router or Server connected to, accessible in DMZ

Setting up the Retain environment Process Steps:

- 1. License Retain.
 - Retain requires a total of 4 licenses to enable Retain BAAS. The first two licenses, base Retain license and the Blackberry module license, are uploaded to the Licenses page in the Retain Server management console. The second two licenses are REST API licenses, REST API base license and the Blackberry REST API license, are installed under the Server Configuration | REST API tab.
- 2. Configure Retain server connection.
 - Found under Server Configuration | Communications tab, the Retain server connection info at the bottom of the page. This connection setting is the information which will be used by the REST applications, (in this case BBM Protected application), or the Retain Router, to connect to the Retain Server. If this is not correct, the Router or the REST applications will not be able to connect the Retain Server. This must be configured before moving on to step

Retain Server Connection

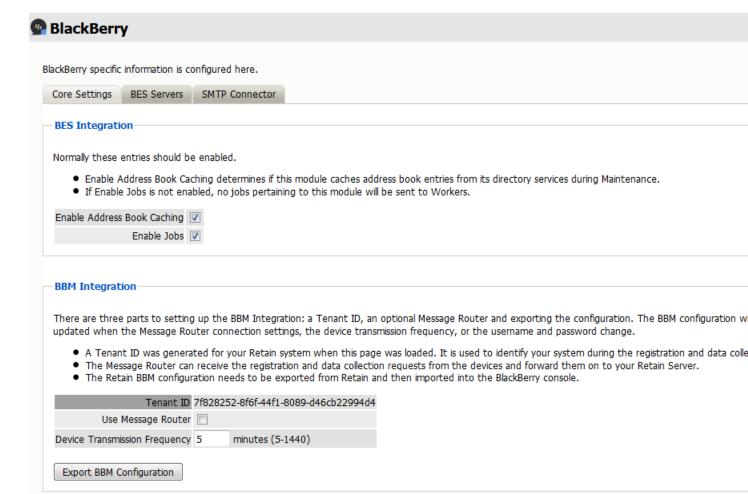
The Router and REST applications will need to know how to connect to the Server. It is highly recommended that you use a DNS host name for the Host Name. It's important to note that you are specifying the URL from the point of view of the **Router or REST application contacting the Server**, which may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



3. From Blackberry module, Export BBM configuration file.
The BBM Configuration file is generated through the Blackberry module configuration page.
Open the module configuration and look at the Core Settings tab. Look at the BBM Integration section. Here the decision is made whether to use the Retain Router or to use the Retain Server.

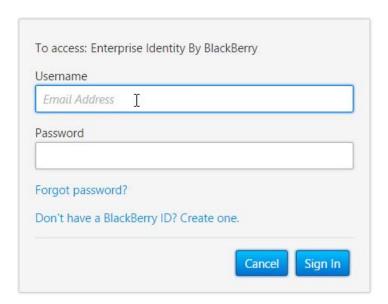
If the Retain Router is to be utilized, select the 'Use Message Router' checkbox.

The Device Transmission Frequency setting is the setting which determines how often the BBM Protected application will upload archive data to the Retain Server. If there is no information to be archived, the BBM Protected application will wait until there is. The setting is in minutes. Minimum is every 5 minutes, maximum is 1440 minutes (24 hours). Once configured, select the 'Export BBM Configuration' button and save the file; you will need it later.



- 4. Obtain a SSL security certificate for destination server/or router (if not already obtained)

 The security certificate must be obtained to configure the Blackberry Identity console. The security certificate needs to be in base 64 format.
- 5. Go to the Blackberry Enterprise Identity Console. (https://idp.blackberry.com/enterprise

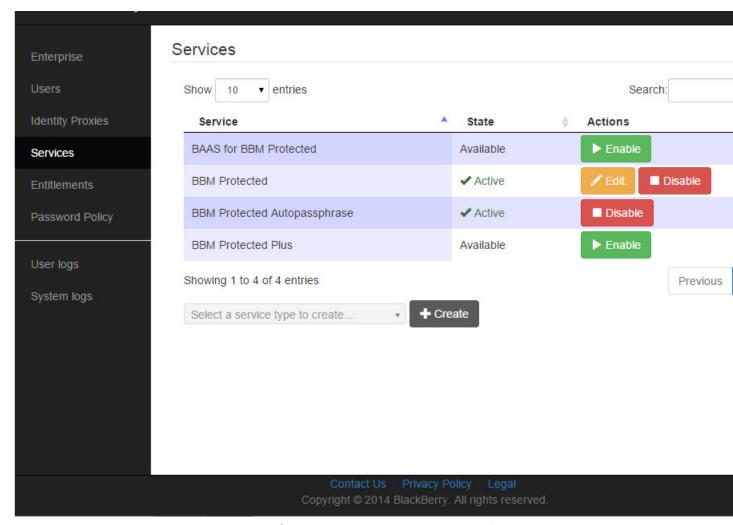


- 6. Enter the credentials provided by Blackberry when account was created.
- 7. Enter administrator console.

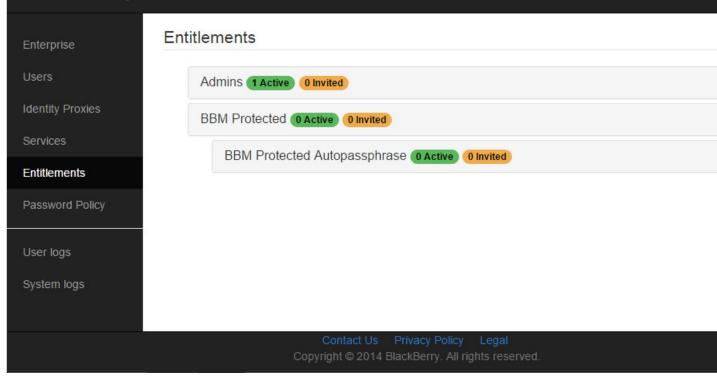
Welcome to Enterprise Identity by BlackBerry. Click the service below that you want to access.



- 8. Select the 'Services' page.
- 9. BBM protected must be enabled.



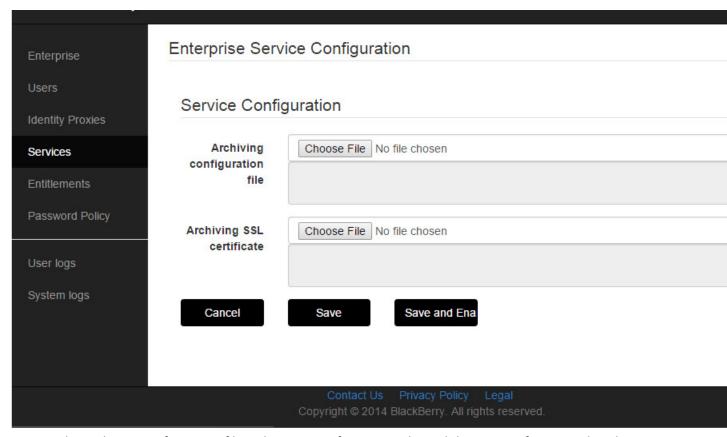
10. Enter 'entitlements' and invite users/devices to be added to the service (In this case BBM Protected)



Users must first be added to the BBM Protected system before they can be added to Retain BAAS.

11. Enable BAAS for Protected.

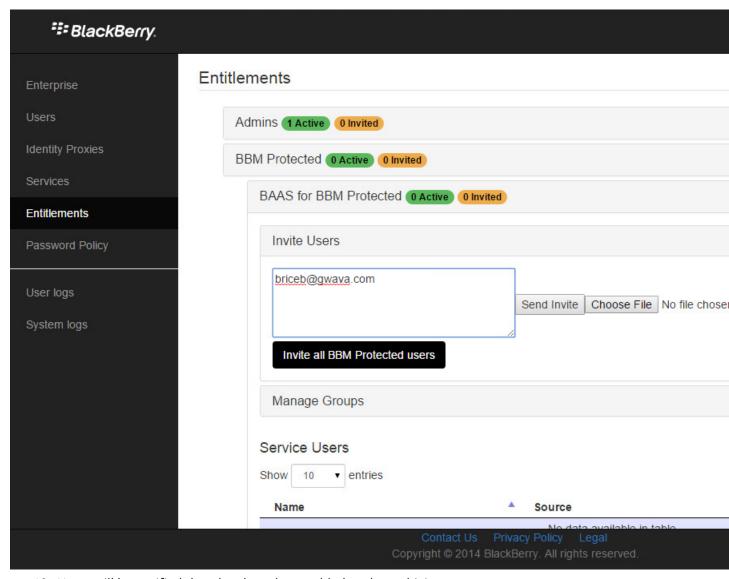
To enable BAAS for protected, from the 'Services' tab, select the 'enable' button for BAAS. There are two files which must be uploaded to the Blackberry Identity Console, these are the two files saved earlier: the SSL security certificate for the Retain Server or the Retain router, (whichever is going to be used), and the BBM Configuration file created in step 3.



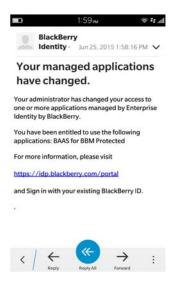
The Archiving configuration file is the BBM Configuration File, and the SSL certificate needs to be uploaded to the 'Archiving SSL Certificate' section. Select the 'Choose file' button and browse to the appropriate file to upload it. Once both files have been uploaded, select the 'Save' button.

12. Enter Entitlements and invite users

Users invited to BAAS MUST already be users of BBM Protected. If users are not registered to use BBM Protected, BAAS will not work. First invite them to BBM Protected, and then invite to BAAS after the device has been registered with the BBM server for BBM Protected service.



13. Users will be notified that they have been added to the archiving



Once the devices and users have been invited to join the BAAS program, they will be sent a notification that they are now part of the BAAS archive.

They are now part of the BAAS system and their BBM Protected communication will be added to the Retain BAAS archive.

Device Mangement

Devices are deleted from Retain through the Device Management interface. Device Management is split into two sections: Mobile and Blackberry. In both sections, multiple devices may be selected and acted on. In addition, all selected devices remain selected until either an action is taken or the administrator has navigated away from the page.

BlackBerry

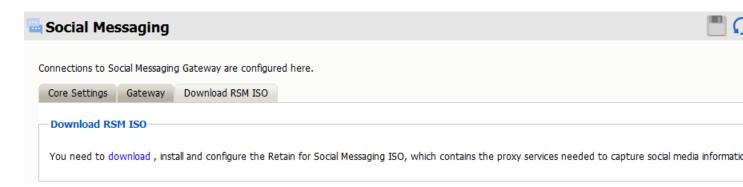
Retain supports archiving data from both BlackBerry devices and devices running the BBM Protected system. Devices running the BBM Protected system can have their communication data archived through the BBM app, obtained from BlackBerry.



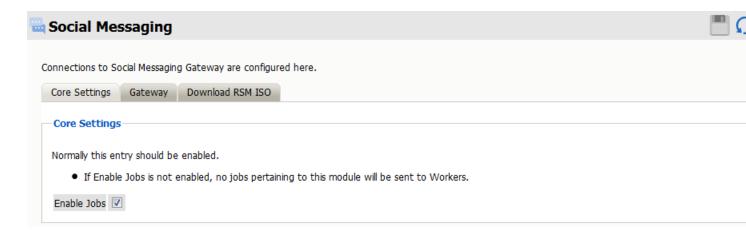
The BlackBerry device page displays all devices which have been registered into the Retain system. Native BlackBerry and BBM Protected devices are displayed together. Retain registers BlackBerry devices through log files from the BES server or through data sent by the device running the BBM Protected application. If a device is not yet registered, it is either not in the BES address book or no data has been archived for the device. The device list may be filtered for any specified content desired. In addition, devices may be selected for removal.

Social Messaging

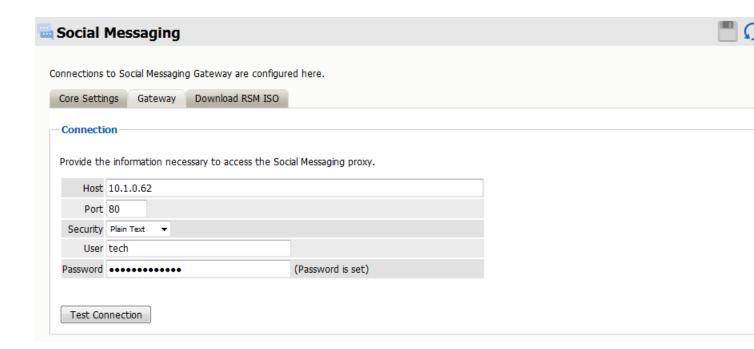
The Social Messaging module allows Retain to archive Social Messaging communications from Twitter, Facebook, Yammer, AOL chats, LinkedIn, Yahoo Messenger, Google Talk, Lync & Skype for Business, and Gmail. The Social Messaging module also archives searches performed on Google, YouTube, Yahoo, Bing, and Wikipeida. To configure the Social Messaging module, all that is required is to provide the connection information and address to the Retain for Social Messaging Gateway. If the RSM Gateway is not yet installed or available, a direct download link is provided on the last tab.



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



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

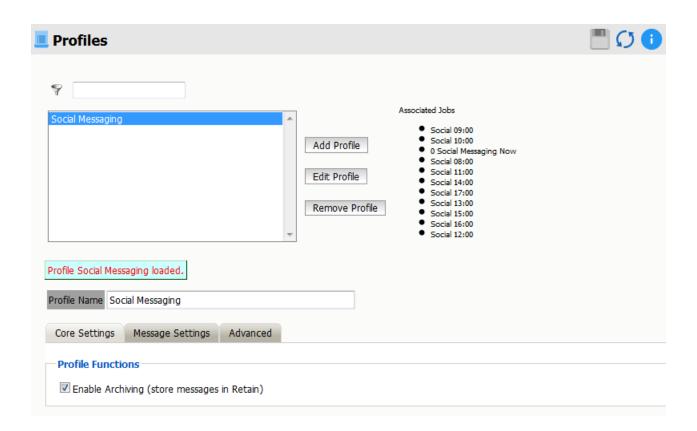


Security and user login credentials are required to connect to the RSM. The user must have administration rights to gather data from the Gateway. By default the user account is 'tech' with the password 'retain'.

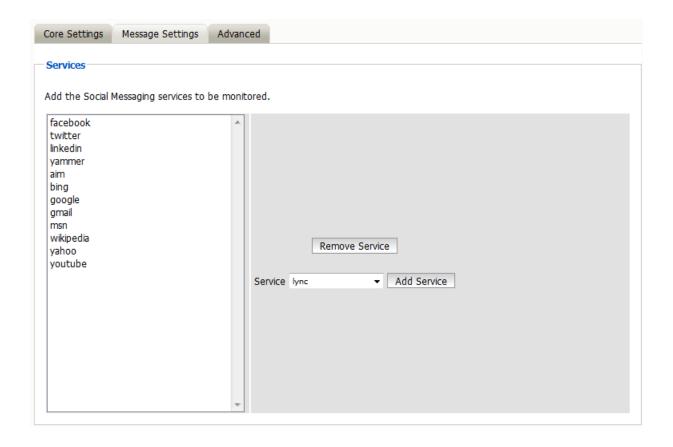
Save changes and the module is configured.

Social Messaging Profile

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



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

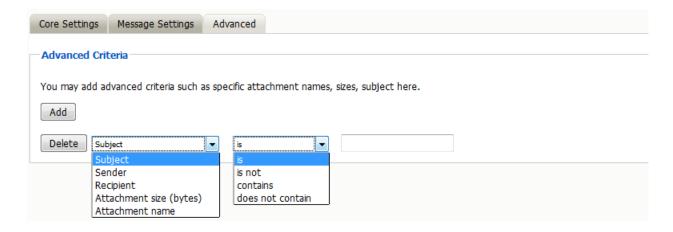


The Message Settings tab can be used to exclude or include specific data types for this profile. By default all types of messages will be archived. If left blank, the default will persist.

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

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

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



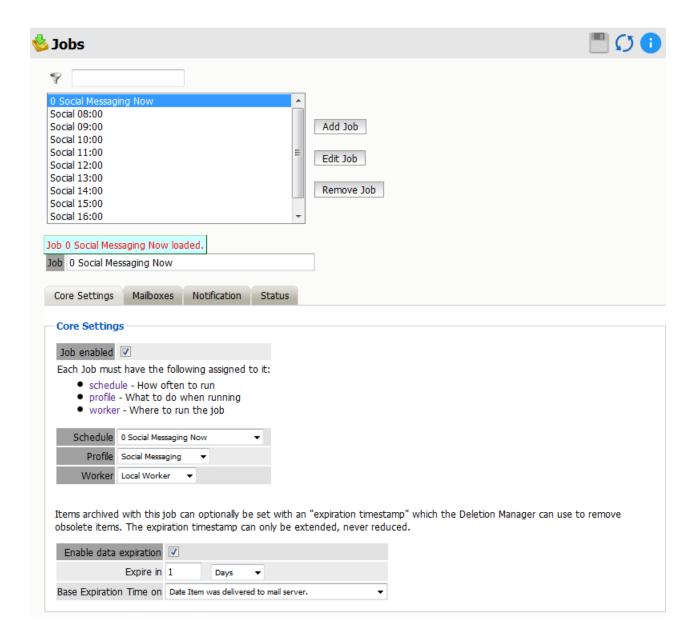
Save changes and the profile is configured.

Social Messaging Jobs

All data collection for the Social Messaging module is configured through the Social Messaging job interface. The jobs here combine the Profile, the Schedule, and a Worker together to archive a specified mailbox, distribution list, or domain. A Job must have a Profile, Schedule, Worker, and a target, mailbox, list, or domain, before it can be saved or run.

Core Settings

The Schedule, Profile, and Worker selected here will determine what is archived, when it is archived, and what worker does the actual archival work. The Job must be enabled in order to run. Only previously configured Schedules, Profiles, and Workers can be selected.



The Job Expiration is an option to place a timestamp on the data in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, Caringo, and Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one (including Retain) may delete the item before its expiration date.

Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for. In order to have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

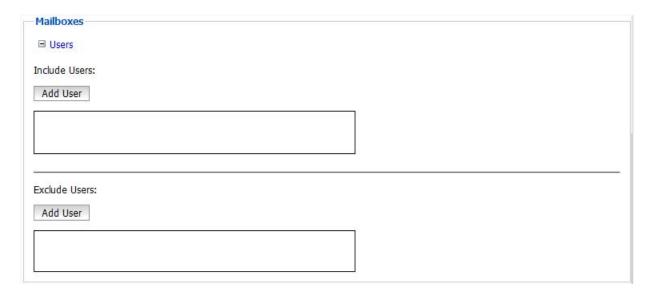


Mailboxes

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

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

If there are users listed, then the users listed in the 'Include Users' box are the ONLY included users, while those listed in the 'Exclude' window are excluded from the job.



Notification

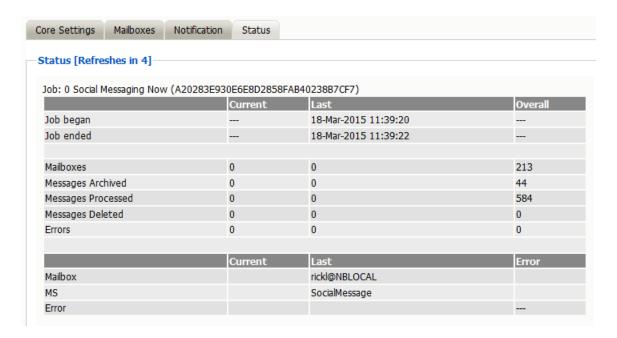
All errors, summaries, and statistics may be sent in an email notification to a desired address on job completion. If it is desired to be notified, configure the notifications under this tab. At least one of the two notification options must be selected in order for an email notification to be created.



Status

On initial creation, there will be nothing of interest shown under the 'Status' tab, as the job is under creation or has not yet run.

However, after a job has been run, or has begun, the 'Status' tab becomes a monitoring tool to report on exactly how much mail and mailboxes have been archived, along with any errors encountered.



Retain for Social Messaging Gateway

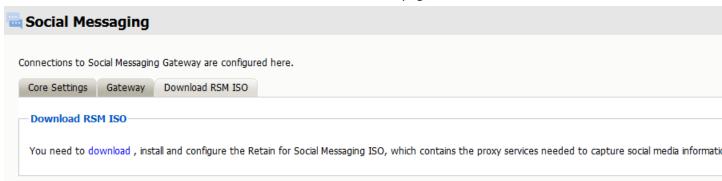
Setting up the VM & Installing the OS

Before setting up the Virtual Machine

- Ensure the firewall is configured as such:
 - o From RSM WAN IP to Untrusted, all TCP/UDP ports
 - o If the RSM WAN IP is a private IP, it needs to be NATed to an appropriate routable IP address. The LAN IP address does not need a corresponding inbound NAT rule.
- NOTE: It is critical that these firewall rules are in place before proceeding.

Configure the Virtual Machine

1. Download the ISO from the link found in the Social module page, on the 'Download RSM ISO' tab.



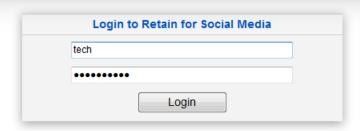
- 2. Install it onto a VM or to the desired machine with the following minimum specs
 - 2GB RAM (minimum)
 - 2 CPU cores (minimum)
 - 60 HDD (minimum)

The recommended RAM, CPU and HDD specification will be determined by the expected load and size of the organization being served. Consult your GWAVA technical representative.

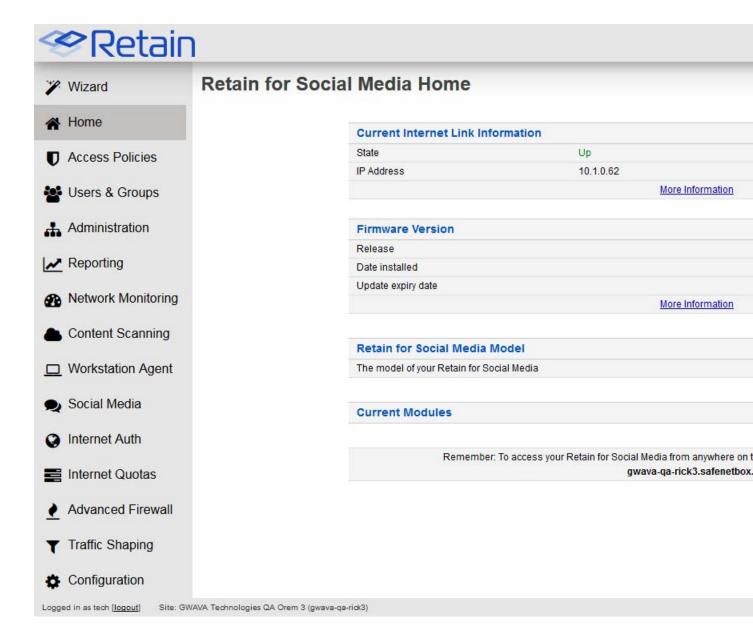
- 2 NICs vmxnet 3, ensuring the interfaces are in different VLANs. NOTE: The second NIC in the VMware settings list is the primary NIC that will be used for proxy traffic. This NIC will be labeled "Internet" on the RSM. The first NIC will be the "Local" port, and may be disconnected if desired when configured.
- Redhat Linux enterprise 5 64bit base OS (selected from the dropdown list when setting up the VM – there is no need to actually install the Redhat OS)
- 3. Mount the ISO and start up the VM or server.
- 4. Set the password for the "tech" and "admin" users. This will then install the RSM OS. Once it's installed and has come back up after the reboot, the RSM will have the default LAN IP of 192.168.0.254/24. The WAN interface will try to obtain an IP address from DHCP.
- 5. Set the LAN IP via the console by sending a ctrl-alt-ins and then follow the prompts.
- 6. All configuration operations are via the web user interface. Log onto the RSM interface by browsing to the LAN IP on either port 80 or 443 and logging in as the "tech" user, (the "admin" user doesn't have the required permissions).







You will initially be prompted to confirm the EULA and provide your details. If the email address is not accepted at this stage (the RSM will try to validate all addresses), then use "root@ mail.RSM".



- 1. Click the "Wizard" button from the left-hand navigation bar, and follow the prompts for setting the language, time zone and network details.
- 2. Once the network settings have been accepted, select
- 3. Configuration ->Internet from the left hand navigation pane, and check "Show Advanced Options".
 - From the drop down list next to "DNS server configuration" select "Always use static server" and then enter in the required DNS servers.
 - Click on Update.
 - Browse to
 - Configuration -> Apply and click on the "Apply Changes" button.

NOTE:

• If the LAN IP Address has been changed in the Wizard setup, you will need to browse to the new IP address and log onto the RSM interface again after applying the network settings.

- The WAN (or Internet) Interface of the RSM will be used for proxy traffic.
- The LAN interface is not used for production traffic, but can be utilized as a management-type port if required.
- It is critical that the two interfaces are configured on different subnets from this point on, regardless if the LAN port is utilized or not.
- 4. After the Network configuration has been completed you will be prompted to supply a sitekey. A sitekey is a unique identifier for your RSM and it should be relevant to your organization. Only alphanumeric and "-"are accepted, with a minimum character count of 3.
 - a. Your RSM Blue representative will supply the registration key to you.
 - b. After the sitekey has been successfully applied, your RSM will reboot.
- 5. Log back onto the RSM and browse to Administration -> Updates and click on "Run Update Now".

The RSM will go through its update process, which could take up to 30 minutes and reboot multiple times.

RSM Configuration

The configuration of the RSM is achieved via its Web User Interface. You need only browse to the (WAN or LAN) IP address and enter in the appropriate credentials. For the following you will need to use the "tech" account with the password that was set when creating the virtual machine. Specific instructions on how to use and configure the different features of the interface are accessed through the 'Help' screens inside RSM.

NOTE: When making changes to the RSM configuration, remember to update the page you are working on before navigating away from it. When you want to make your changes live, browse to the "Apply" section in the relevant module and click on the "Apply Changes" button.

Proxy Configuration

The following describes how to configure the RSM in a direct proxy setup. The client will then direct their browser to the WAN IP address of the RSM on the specified port (8080 for example).

- 1. Browse to Configuration -> Web Proxy.
- 2. In the drop down list next to "Direct Proxy Mode" select "Direct".
- 3. Specify the port to use (8080 by default).
- 4. In the drop down list next to "Provide proxy on Internet interface", select "yes".
- 5. In the drop down list next to "HTTPS inspection", select "Enabled for all traffic".
- 6. Click on "Update".
- 7. Browse to Configuration -> Apply and click on the "Apply Changes" button.
- 8. After 1-2 minutes, once the changes have been applied (when the yellow apply banner disappears), test the proxy by pointing your browser to the RSM WAN IP on the specified port and browse to a standard HTTP web site (HTTPS sites will be tested shortly).NOTE: You can view the proxy logs by browsing to Configuration -> Web Proxy and clicking the "View Web Proxy logs" link in the top right hand corner of the page.

RSM CA Certificate Installation

The RSM will be performing HTTPS inspection, meaning it will have visibility to view encrypted web traffic. This is critical to ensure that all Social Information Governance functionality is available. In order to do this successfully, the RSM CA Certificate needs to be installed on all client devices that will be using the RSM as a proxy service.

- 1. Download the certificate from <IP>/noauth/cacert.
 - If using Microsoft Active Directory create a Group Policy Object
 - In GPMC open the relevant policies that apply to the Computers that need to have this certificate installed and navigate to Computer Config>Windows Settings>Security Settings>Public Key Policies and import the CA Certificate as a Trusted Root Authority.

NOTE: This GPO will NOT apply to Safari or Firefox browsers. The certificate will need to be manually installed.

- 2. If manually installing the certificate ensure it's installed as a Trusted Root certificate.
- Once the certificate has been deployed, test by browsing to a HTTPS site and view the web proxy logs. You should not see a certificate warning in the browser and the proxy logs should display your HTTPS request.

Social Information Governance Configuration

Social Information Governance is a Rule/ Action process, meaning that a rule needs to be triggered before an action takes place. A rule is made up of one or more criteria. A criterion tells the RSM what should trigger the rule. A rule has one action assigned to it. An action may have multiple sub-actions. Actions can be either:

- Block
- Alert
- Modify
- Moderate
- Or in some cases, combinations of the above (i.e., Block and Alert) For example a rule to block a staff member from posting extreme profanity on twitter will contain the following:
 - o A criterion that uses the "extreme profanity" pattern list.
 - o Another criterion that sets the Application type to "Twitter".
 - o A final criterion that specifies the user's action as "send"
 - o An action to block.

NOTE: All Criteria must be true before the rule is triggered. In other words: If the user sends a post to Twitter, which matches an entry in the "extreme profanity" pattern list, then take the associated action to Block. This rule would not trigger if the user read extreme profanity on Twitter, so care must be taken in establishing the Social Information Governance requirements for your organization. The RSM comes with predefined rules that may apply to the relevant organization type.

Enabling Social Information Governance

- 1. Enable Social Information Governance by browsing to Content Scanning -> General. Under the SafeChat Settings table:
 - a. From the dropdown list next to "Enable SafeChat scanning of web content?" select "Yes"
- 2. Click on "Update"
- 3. Browse to Content Scanning -> Apply and click on the "Apply Changes' button.

Creating Pre-defined Social Information Governance Rules and Actions

- 1. After the changes have been applied create the predefined rules by browsing to Content Scanning -> General and click on the "Add suggested settings" button.
- 2. Browse to Content Scanning -> Rules where you can see the rules that have been created. The rules and corresponding actions will be currently disabled.
- 3. Enable the rules by
 - Clicking "edit" next to the rule.
 - Check "Enabled".
 - Additional criteria can be added at this stage by clicking on the "Add Criteria" link.
 - Click on "Update".
- 4. When you are happy for the rules to come into effect, browse to Content Scanning -> Apply and click the "Apply Changes" button.

Creating Custom Social Information Governance Rules

To add or modify a rule:

- 1. Browse to Content Scanning -> Rules and click on "Add new rule" in the top left hand corner
- 2. Enter a descriptive name for the rule into the textbox provided. This will be used for referring to this rule, for example when modifying or deleting this rule.
- 3. Select the action to be run when the rule is matched.
- 4. Click the Update button to save the rule.
- 5. Enter criteria as required. The rule will be triggered when all of the listed criteria are matched. NOTE: If no criteria are specified then this rule will never be triggered.
- 6. Select the checkbox provided at the top of the page to enable the rule, once you are happy for it to become active.
- 7. Click the Update button to save the rule.

Creating Custom Social Information Governance Actions

To add or modify an action:

- 1. Enter a descriptive name for the action into the textbox provided. This will be used for referring to this action, for example when modifying, deleting or using this action.
- 2. Click the Update button to save the action.
- 3. Enter sub-actions as required.
 - NOTE: If no sub-actions are specified then this action will not do anything, but the rule will still trigger a match, this can be useful for reporting.
- 4. Select the checkbox provided at the top of the page to enable the action, once you are happy for it to become active.
- 5. Click the Update button to save the action. Once all changes have been made to the
- 6. Content Scanning section:
 - a. Browse to Content Scanning -> Apply.
 - b. Click on the "Apply Changes" button.

The Configured Social Information Governance rules and actions will now be live in a few seconds.

Social Information Governance Moderation

The Moderation section of Content Scanning allows authorized users to approve or deny messages held by a "Hold for Moderation" sub-action. A Content Scanning action needs to be configured to hold messages in a particular moderation "queue" until processed. An authorized user can then view each queue, reviewing the pending messages of a queue in detail, and selecting whether each message is approved or denied. An example would be to hold for moderation any messages that contain the name of the company for moderation and review before it's sent to Twitter, Facebook, etc. The authorized user can then elect to allow that message to be sent or reject the message, giving an explanatory comment if necessary.

NOTE: Please refer to the online help under Users & Groups to create individual users and groups if you want to limit access to these features a defined group of users. Social Information Governance Moderation requires the configuration of:

- 1. A moderation queue and;
- 2. An action to hold messages in the moderation queue.

Creating a Moderation Queue

- 1. Browse to Content Scanning -> Manage Moderation Queues.
- 2. Select "Create New Queue".
 - Provide a descriptive name for the queue.
 - Select which groups (if any) can view this queue.
 - If an explanatory comment is required for any actions within this queue (i.e., releasing or denying), set the appropriate value here.
- 4. Click on "Update".

Creating a Moderation Action

This is simply a Content Scanning action to hold messages into the specified moderation queue.

- 1. Follow the steps described under "Creating Custom Social Information Governance Actions".
- 2. When selecting the sub-action, select "Hold for Moderation".
- 3. Select the moderation queue from the drop-down list.
- 4. Assign this action to the required rule (please refer to "Creating Custom Social Information Governance rules" for more detail).

Configuring Secure Social Media Authentication

The RSM can be used to grant users access to social media accounts without having to divulge the account password. The RSM will determine whether users can access particular social media accounts based on the group membership of the user. This can be used, for example, to allow certain members of your organization the ability to make posts to your organization's Twitter or Facebook account without divulging the password for those accounts.

NOTE: This functionality currently applies to LinkedIn, Twitter and Facebook accounts.

Configuring RSM to authenticate to a company's social media site

- 1. Browse to Social Media -> Accounts and click on "Configure New Account" in the top right hand corner of the page.
- 2. Configure the account options:
 - Select the application.

- Type in the username and password for the account NOTE: This feature will not work with accounts that require two-factor authentication. It is recommended that a long, complex and very secure password be used for the account.
- Select which user groups can access this account through the RSM.
 NOTE: Users with Social Media Account administration privileges are allowed to access all social media accounts regardless of this setting.
- Provide a description of the account (i.e., "ABC Corp. Main Facebook page")
- Set a session lifetime. This determines how long the login session cookie will be active before timing out. Leave it blank to use the application's default values.
- 3. Click on "Update".
- 4. Browse to
- 5. Content Scanning -> Apply and click on the "Apply Changes" button.
- 6. Users can access the company's social media sites by browsing to the RSM, supplying appropriate credentials and then browse to
- 7. Social Media -> Accounts.
 - They will see the list of social media accounts they have permission to access (for example Facebook and Twitter, but not LinkedIn).

Click on the social media account desired, and the RSM will log into that application for the user.

CellTrust SecureLine

The CellTrust SecureLine module allows the collection of phone data from the CellTrust Secureline server into the Retain system. CellTrust Secureline supports the archival of sms and phone call data from both Android and iOS phones. All devices which are supported by the Secureline server are covered by Retain.

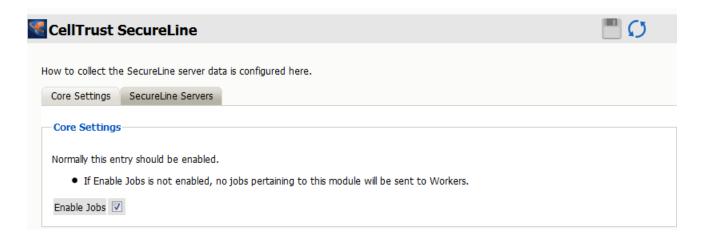
To archive CellTrust SecureLine information, the SecureLine server must be configured to log phone calls, PIN and SMS data. The server log files are what Retain uses to collect the information about phone activity. If the SecureLine server has not been configured to log the data, Retain will be unable to archive the information. See appendix J for Secureline server configuration.

Configuring the CellTrust SecureLine Module is incredibly simple. Essentially, the module needs to be enabled and configured with the log file location. Afterwards, the data is treated the same as any other data entering the Retain archive. Data is archived in the database under the user/phone it originated from and can be searched and viewed through the Search Messages interface.

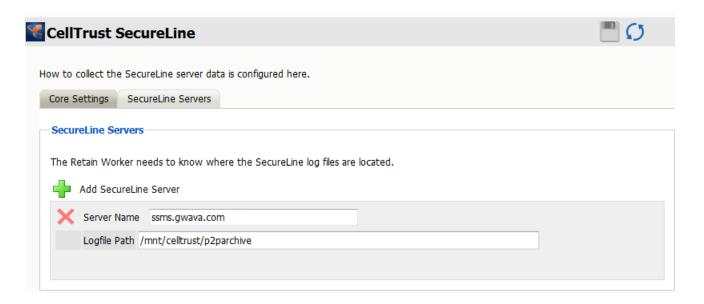


To enable the CellTrust module, select the 'Configure' link under the 'CellTrust' icon from the module page.

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



The log file path for each server in the system is required to archive the phone data. To add a server to the Retain process, select the green plus sign and then configure the server name, (any name will do), and then the log file path.



The log files must be directly accessible to the worker tasked with archiving the CellTrust SecureLine data, whether through direct access or through a mapped drive. The logfile path is the direct file access path to the log files of the CellTrust Secureline server, (/opt/PM/scheduler/external by default), and should be specified as the worker will access the log files.

Archiving/Data Collection - About Workers, Profiles, Schedules, and Jobs

Getting data into the archive is the core functions of Retain. As such, sub-functions have been broken out to give you as much flexibility as possible in deciding your archive policy.

Your archiving policy will cover several points:

- What to archive (jobs)
- Includes what message server(s) to collect data from.
- What worker
- When will the data be collected (schedule)
- How to collect the data will cover (profile)
- Types of items or users to archive.
- Date range of items to be archived.
- Avoiding redundant data collection.
- Utilizing storage flags to achieve 100% retention.
- Whether to store attachments.
- Advanced criteria, such as attachment names, or folder locations.

Retain Job components

Schedule

The Schedule is the time framework that is set for the job to run on and provides the choice of how often, and when, the job is run. Single time, (job runs once), or a recurring time frame are available. Schedules are universal and are not module or job specific. They may be used for one or all jobs.

Profile

The Profile dictates which kinds of files to retrieve from the message system. All message system items are selectable to be archived, but you may also specify time periods, exclude folders or subjects, and how the data is stored. Profiles are module specific, and the profile section contains sections specified by each module and information only applicable to that module.

The Profile also has a very important function, duplicate checking. This keeps the Retain system from storing multiple instances of the same message. Make sure you check and read the section on Purge (GroupWise), Retention (GroupWise), and Item Store (Exchange/GroupWise) flags.

Worker

The worker is the part of Retain that actually pulls the data from the message system servers. The Worker contacts the Retain Server and gathers information on which jobs are associated with it, when to start them, and which items to get. Workers can only run one job at a time. Workers are not module specific, but may require module specific settings. These settings, where applicable, are found under the module specific tab.

Jobs

The Jobs section is where the Worker, Profile, and Schedule are all brought together and associated into a comprehensive configuration by allowing the user to tie their configuration to a selected Secureline server and selected user list. The desired SeucureLine server must be selected. The Worker will not start archiving anything until a created job is bound to a Worker, Schedule, and Profile.

Google Apps Module

The Google Apps module allows Retain to archive Gmail data items. To configure Retain for Gmail archiving, Retain needs Gmail to be configured to allow Retain access, and the appropriate information entered into Retain.

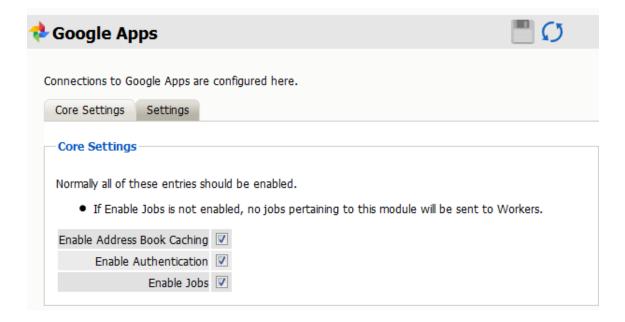
Google Apps requires that a project be created, an OAuth key created and a Service Account specified and enabled before the Retain system can connect and archive mail.

To archive from Google Apps:

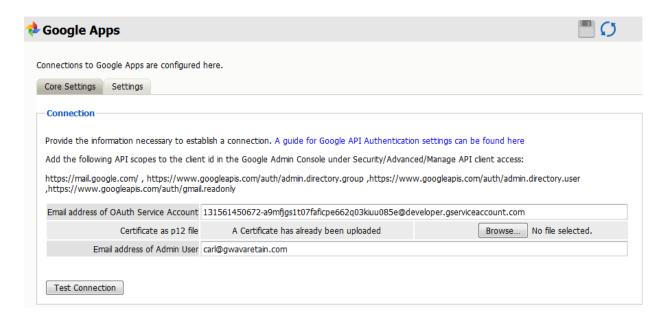
- Create a project
- Create a new Client ID key.
- Generate a P12 key certificate

For updated information on configuring Google Apps, and the Google Apps module, see: http://sup-port.gwava.com/kb/?View=entry&EntryID=2516.

Once Google has been configured to allow access to Retain, the Google Apps Module may be configured.



Address book caching must be enabled to gather and maintain an updated list of users. Authentication is used to allow access to the Retain message store for users based on their existing Gmail account login. If the Enable Jobs option is not enabled, no jobs may be completed with the Google Apps module.



In order to connect to the Google system, Retain requires the email address of the Admin user OAuth Service Account and a p12 Certificate for authentication. Retain archives the Gmail system through IMAP, and will login and download the message data to the Retain data store.

Test the connection to ensure that the configuration has been completed correctly.

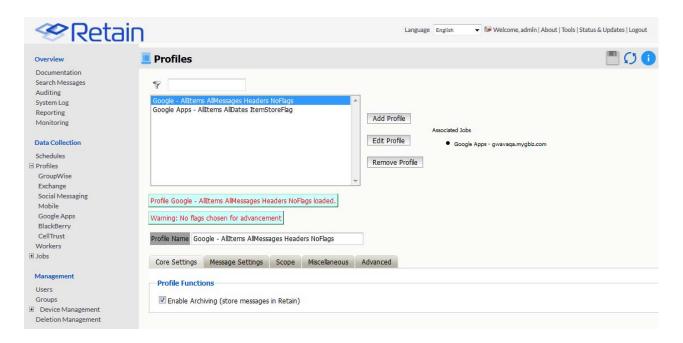
Jobs and profiles for Gmail will not be visible until the address book has been cached. After the module has been cached, all configuration options for profiles, workers, schedules, jobs and data storage will be enabled and visible.

Google Apps Profile

To create a profile to archive email and data from Google Apps, select the add profiles button and name the profile then continue configuring as desired. The Google Apps profile must be configured and all settings saved before a job can be created and run.

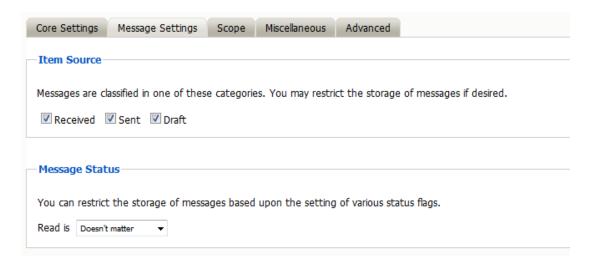
Core Settings

The profile will not become active and will not allow jobs to be run unless the profile is enabled. Enable the profile by placing a check mark in the 'Enable Archiving' checkbox.



Message Settings

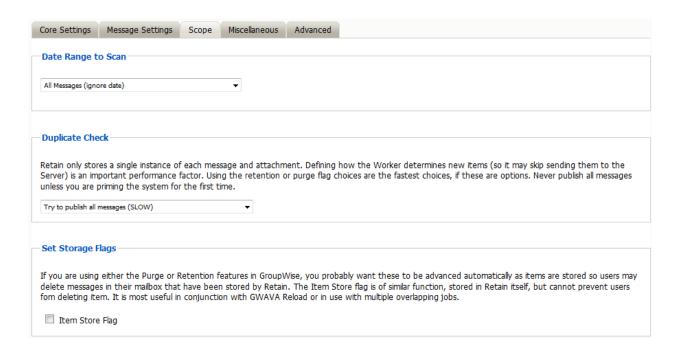
The Message Settings tab contains the source and status settings for the messages to be archived. Message types and sources which are checked will be archived in this profile.



Any item type or status which is not selected or specified will exclude qualifying items from being archived in Retain.

Scope

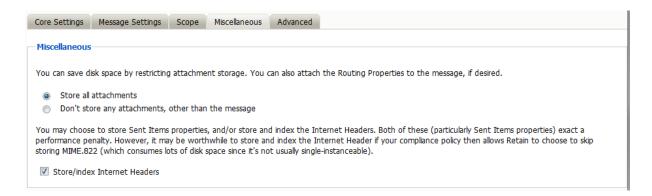
The Scope tab covers the Date Range, Duplicate Check, and Storage Flags. The date range limits the data items archived based on when it was created or sent. The Duplicate check is the method used to eliminate already archived messages from an archival job. Initial jobs should be left to archive all messages, as there are no messages currently archived in the system. Afterwards, the duplicate check should be set to use the storage flag which only attempts to archive messages which are not already in the archive, thus speeding the archival process.



Storage Flags are markers which Retain uses to keep duplicate data items from being added into the archive. This option should be enabled to allow Retain to only archive original data.

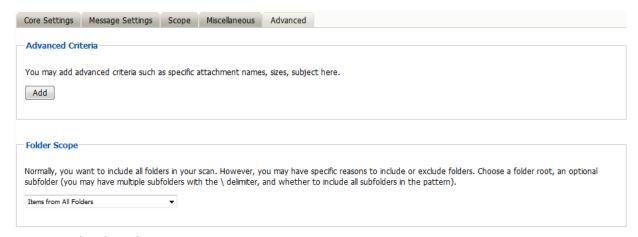
Miscellaneous

The miscellaneous tab allows the configuration to allow or deny archiving and indexing of attachments in Retain. If attachments are to be archived, they may also be indexed to provide searching capability in the browse messages interface.



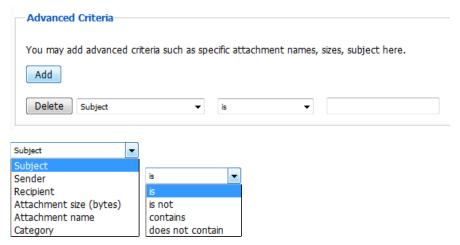
Advanced

The advanced tab allows the administrator to be even more specific in what to archive and what to exclude. The criteria added under the 'advanced' tab will limit associated jobs to only the items matching the criteria. Up to 6 lines of advanced criteria may be added to each profile. Each additional line will be logically AND-ed together. For example, the system will archive all items where the following is true: Criteria A and Criteria B and Criteria C etc.



You may select based on:

- Subject
- Sender
- Recipient
- Size
- Attachment Name



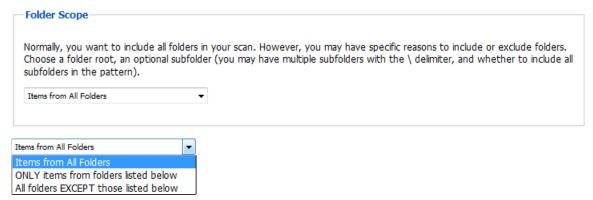
Specify the criteria, whether it is: equal to or not equal to, whether they contain or do not contain the item.

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

Folder Scope

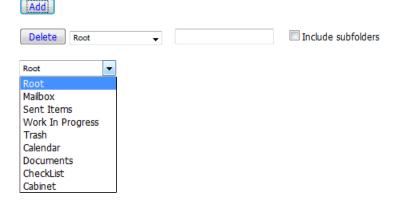
By default we dredge items from all folders. You can specify one or more inclusions or exclusions. Your choices are:

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



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

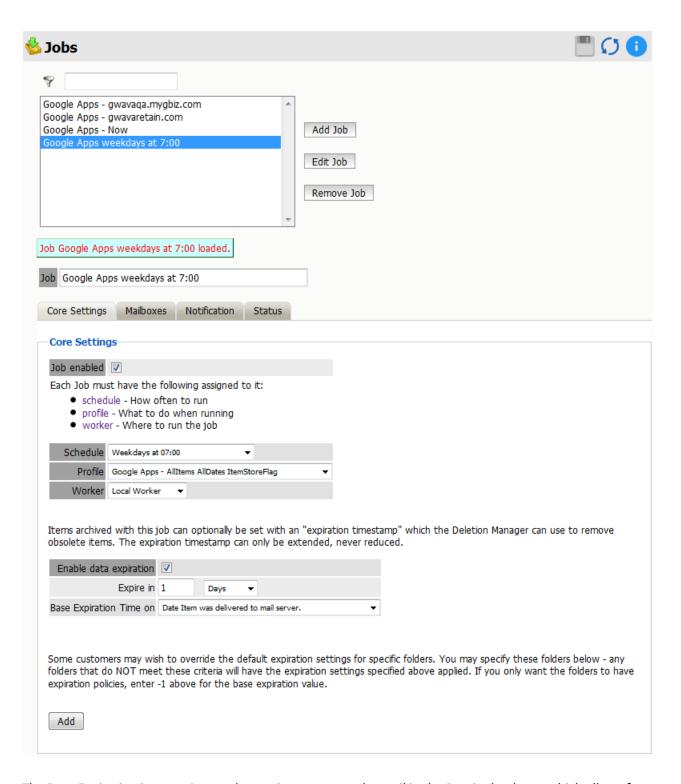
- 1. Select to only get items from the list, or all except items from the list, as desired
- 2. Select 'Add' to open a new selection
- 3. Specify a System Folder (mandatory). Example: Calendar
- 4. You may specify a subfolder of that folder (optional) Example: entering "old" would mean the folder "old" under "Calendar"
- You can have multiple hierarchies under that with the / delimiterExample "old/mail" would mean the subfolder "mail" under "Old" under "Calendar"
- 6. You specify if the option includes subfolder Example: If you select "old" and "includes subfolder" is unchecked, "Calendar/mail" is selected. If "includes subfolder" is checked, "Calendar/old/mail" would also be selected



Google Apps Jobs

All data collection for Google Apps is configured through the Google Apps job interface. The jobs here combine the Profile, the Schedule, and a Worker together to archive a specified mailbox, distribution list, or domain. A Job must have a Profile, Schedule, Worker, and a target, mailbox, list, or domain, before it can be saved or run.

The Schedule, Profile, and Worker selected here will determine what is archived, when it is archived, and what worker does the actual archival work. The Job MUST be enabled in order to run. Only previously configured Schedules, Profiles, and Workers can be selected.



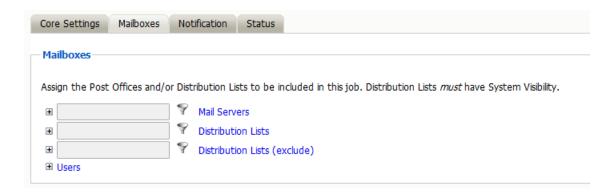
The Data Expiration is an option to place a timestamp on the mail in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, Caringo, and Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one (including Retain) may delete the item before its expiration date.

Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for. The base folders and criteria are specified under the custom expiration dates accessed through the 'Add' button at the bottom. All messages included in any specified folders will have a different date or be exempted from the standard expiration date. In order to have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

Mailboxes

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

Expand the Post Office and/or Distribution List trees, and check off the items you want to be dredged. Note: If you desire to have a job backup a single user, or selected group of users, select the Users menu and assign the user(s) desired



The users section allows you to select individual users to include, or exclude them from an archive job. For example: you can select an entire Mail Server to be archived, and then expand the users section to include or exclude users to the job.

This can also be used to select only certain users in the system for an archive job.

To add a user to the Include or Exclude list, select the respective 'Add user' button and search for the user. It can be helpful to unselect the 'only show recently cached items' option.

Add the selected users to the list in the search window, then select 'Ok' to add them to the include or exclude list.

Notification

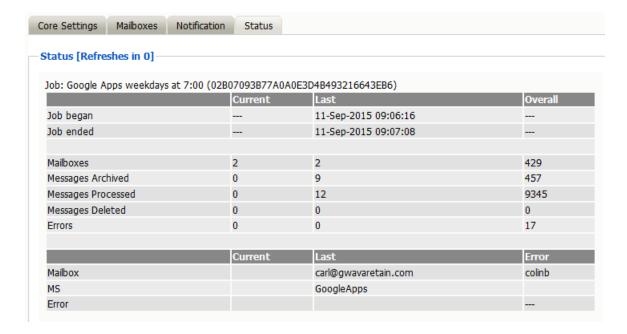
When a job is run, the notification option allows the administrator to be emailed a summary of each running job if desired.



For notification to function correctly, the SMTP information for the desired SMTP server must be fully filled-out. How much information is required varied depending on the mail system used.

Status

The Status tab displays the status of any currently running jobs, as well as the stats of the last completed job.



This tab is informational only unless a job is currently running. If the selected job is running, an option to abort the job is displayed.

This window will refresh every ten seconds to keep you up to date as to the status of the selected job. The completed line is displayed during an active job. The completed status is a display of how many mail-boxes have been completed, the job mailbox total, and gives an incrementing percentage for the amount completed. This amount is based entirely on the number of mailboxes, not the amount of mail. Because the last mailbox could be larger than the rest of the system, this percentage may not be accurate according to time.

Mobile Module

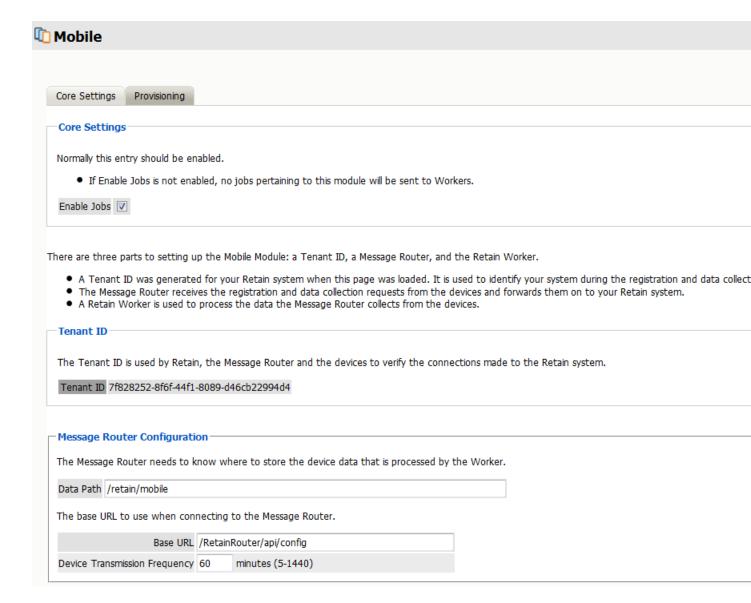
The Retain for Mobile module allows the capture and archiving of MMS, SMS, and phone call information from Android smart phones. (Android 4.0.4+ is supported.) To enable mobile data archiving the following tasks must be completed, in order:

- 1. The Router must be installed and be available outside the host network. (in DMZ)
- 2. The Mobile Module must be configured, and a Mobile worker created.
- 3. The Retain Android application must be installed to users' phones.
- 4. Devices must be registered under 'Device Management'.
- Users must use the registration email to verify device registration with the GWAVA server.
- 6. The data path folder(s) must be manually created.

The Retain Router should be installed in the DMZ area of the host network. The Router must be available to the Internet and the Retain Server. The Router and Worker may be installed on the same machine, or apart, however both the Router and the Worker must have access to the Data Path location. If they are not to be installed on the same machine, a mapped drive must be created to the location. See the install guides for more information.

The Retain.apk application can be downloaded from the tools page of the Retain Server. It is recommended to use a Mobile Device Management (MDM) system to distribute the APK to devices. When using an MDM to distribute the Retain app, the app can be set to not allow app removal. The MDM will need the package name for the Retain app which should be set to com.gwava.retain.mobile. Consult the MDM's documentation for instructions.

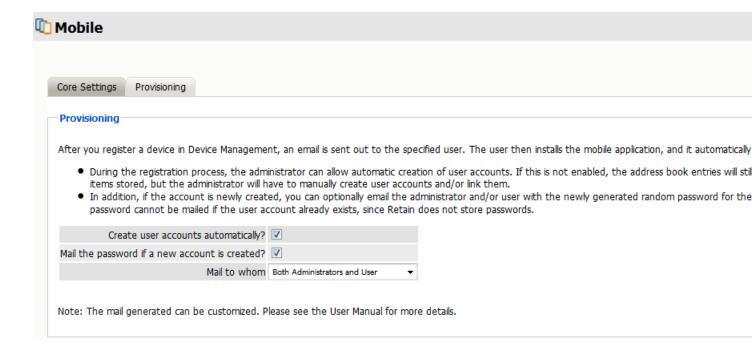
The application sends collected information, approximately every hour, to the Retain Router. No data will be archived until the application is registered and configured.



Data Path

The data path is the store location for messages received by the router, but not yet sent to the Retain Server. This must be a location which the router will have constant and reliable access to. This location must also be accessible to the worker assigned to the router. A mapped drive will be required if the data path is not located on the same box as the mobile worker and mobile router. This path should be the path to the data directory from the Router, and should be absolute. Folders will not be created automatically. If the data path is being created on a Linux system, the ownership should be tomcat:tomcat. After entering the configuration, test the connection.

Provisioning options allow for automatic user accounts, and associated passwords mailed to the designated address.



After the module configuration has been saved, a worker must be created to allow the information to be accepted into the Retain Server. The Mobile worker, Profile, and Device Management menu pages will not be available until the Mobile Module has been configured, and the configuration has been saved.

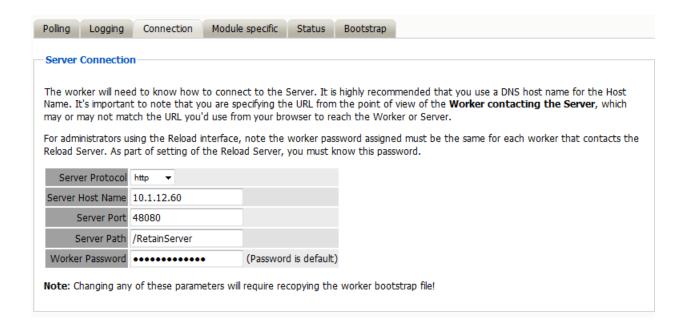
Worker

The Mobile worker is a special worker configuration which contacts the Retain Router with all the information it needs to function. This includes the configuration settings for the data path and connection information. In general, the Mobile Worker is configured the exact same as most workers in the Retain system, however there is one setting for the Mobile Worker which is specific to the Mobile Module; the data path. The data path is essential for the worker to function with the Mobile Module. An existing worker may be used with the Mobile Module, however, it must have the bootstrap file reuploaded once the data path has been configured.

The mobile data is delivered to the Retain Server through a router service, which must be informed how to contact the Retain Server. While the Router is gathering information between jobs, the mobile data is stored. The stored location is specified by the data path. This is why the Router and the worker must have direct and constant access to the data path; the data path stores the message data collected by the Router and retrieved by the Worker.

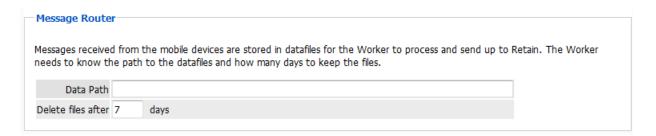
To ensure connection to both the data path and the Retain Server are correct, the Connection and the Module Specific tabs must be reviewed and configured. The Connection tab holds the connection address which the Retain Worker will use to contact the Retain Server. Depending on where the Server and the Worker are located in the network, and because the Worker may be installed alongside the Router in the DMZ, the connection information must be accurate to allow a clean connection to the Retain Server. An IP Address will work, but the DNS hostname of the Retain Server is also recognized and supported.

The port and address must be open or forwarded through security and firewalls from the Worker to the Server. In addition, the Retain Server will not accept connections without the password set here. The password is randomized and doesn't need to be changed.



Set the correct information and move to the 'Module Specific' tab.

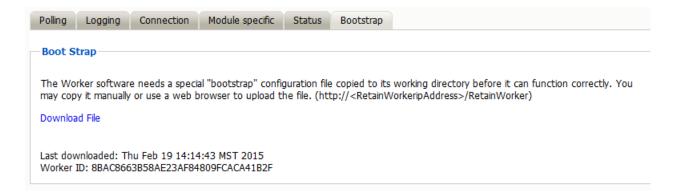
Under the Module Specific tab, the only field that applies to the Mobile Module is the field at the bottom marked 'Message Router'.



The path that the Worker will use to connect to the mobile Data directory should be placed here. This may be a mapped drive or a local path. The path should be absolute. The connection to the data directory is essential to the worker function, if this is unavailable or blocked, mobile jobs will fail.

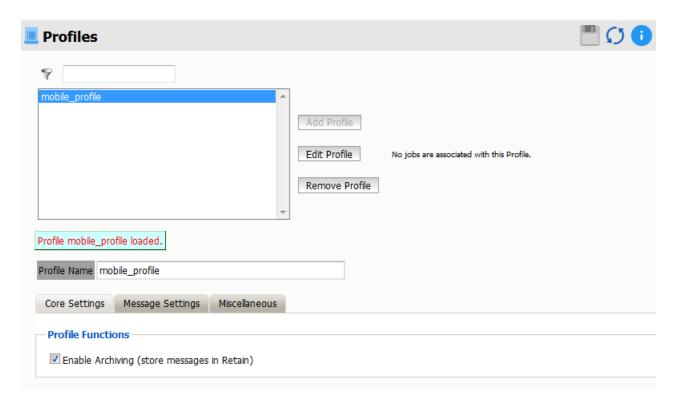
The worker will delete files after adding them to Retain according to the time frame listed. A setting of '0' will tell the worker to delete files immediately after importing them into the Retain Server.

Once the Mobile Worker has been configured, save the changes. Once the settings have been saved, the worker will be created, and the configuration saved to the bootstrap file. The bootstrap file must be uploaded to the Retain Worker before the Worker will function. Select the 'Download File' link to save the configuration file. Browse to the worker's page, and upload the configuration file. Once uploaded, the worker automatically reads the configuration and checks with the Retain Server for jobs.



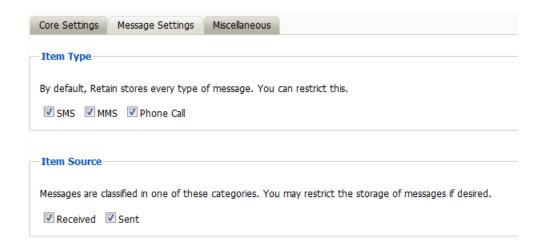
Profile

The Mobile Profile allows the administrator to dictate what types of message data is collected from registered mobile devices. The profile is universal, and once configured works for the entire mobile system.

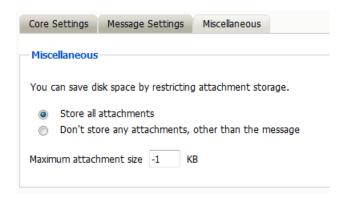


The Mobile Profile is fairly straight-forward. To enable message data collection for the mobile system, the profile must be activated.

Under Message Settings, the different types of messages which can be archived from Mobile devices are listed and configured. The message source and type are available for selection or exclusion.



Finally, the Miscellaneous tab allows for configuring attachments. Attachment size and general attachment settings can be configured here. If there is a maximum limit, set it in the field below.



The attachment size limit is listed in KB. If the size limit is configured to -1, then there is no limit and all attachments will be archived no matter the size.

There can only be one profile for the Mobile module. Once a profile has been created, the option to create a new profile will be disabled and grayed-out. If a new profile is desired, the existing profile must first be deleted. If there is no profile active in the system, the option to create a new profile will be active.

Once the module, worker, and profile have been configured, devices must be added through Device Management.

Device Management

Devices are added to and deleted from Retain through the Device Management interface. Device Management is split into two sections: Mobile and Blackberry. In both sections, multiple devices may be selected and acted on. In addition, all selected devices remain selected until either an action is taken or the administrator has navigated away from the page.

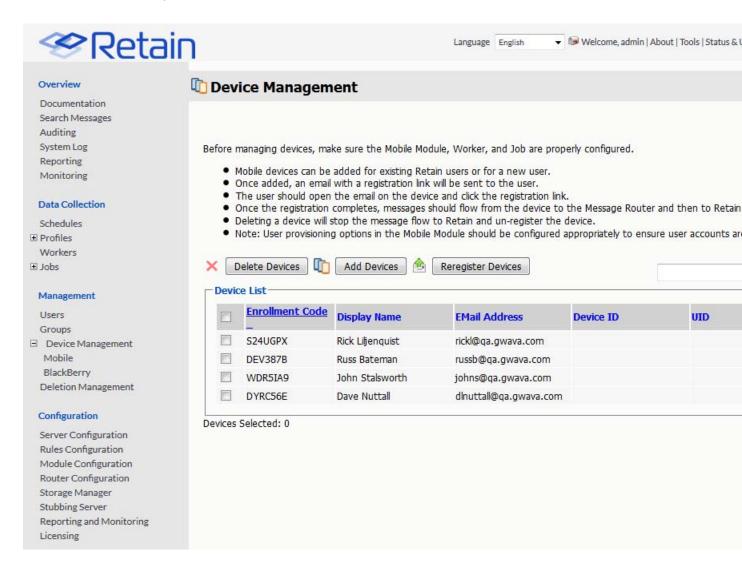
Mobile

Once a device is added to the interface, Retain creates an email which contains an activating link. Users need to open this email and follow the link on their mobile device. The link should be activated through the Retain application on the Android device. The activation email can be customized. The template this email is generated from is located in the file:

...\RetainServer\WEB-INF\classes\config\mobile\registeruser.html

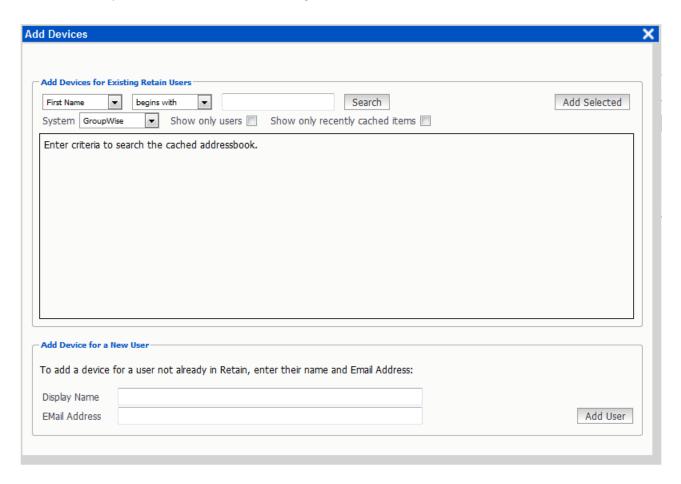
The email sent to users may be customized, including any specific instructions for the organization, can be added or modified here, and all emails sent to users will contain information placed in this template. Information on the email variables is listed at the end of this page.

When a device is added, an email is sent to he user with a unique enrollment code. Devices which have been activated will show a Device ID and a UID. Once registered, the Retain app will start collecting data setup in the profile section, (SMS, MMS, and phone call), and will send it to the Mobile Router to be archived in the Retain System.



To add a device to the Retain system, open the Device Management window and select the 'Add Devices' button.

Devices are only enrolled once. If a user acquires a new phone, or if they change phones, a new registration needs to be completed to maintain the correct codes. To re-register a device, select the desired device and press the "Register Devices" button. Once pressed, Retain will recreate the enrollment, remove the old Device ID, and resend the registration email. The end user will need to repeat the enrollment process stated in the email to register the new device.



The Add Devices window opens, which allows administrators to select and add multiple users. The system may also add users which are not cached in the Retain system, but they must be added one at a time. Manually input the name and associated email address at the bottom and select the 'add user' button.

Once devices have been added and registered, they are accessible to be modified, removed from, and archived by the Retain system.

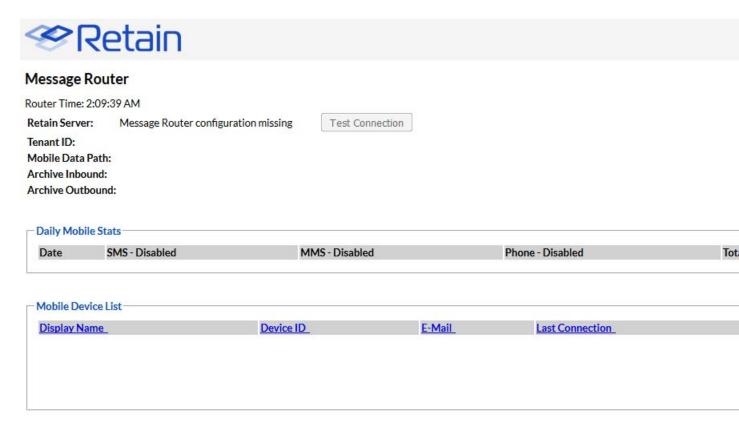
Mobile Router Page

The Retain Router has its own web page. This page can be used to test the connection to the Retain Server, as well as verify that communication is open and that devices registered in the Retain Server are listed in the Retain Router.

To access the Router page, open a browser and enter the connection URL.

http://Retain_Router_IP_or_Domain/RetainRouter

For example: http://192.168.1.21/RetainRouter

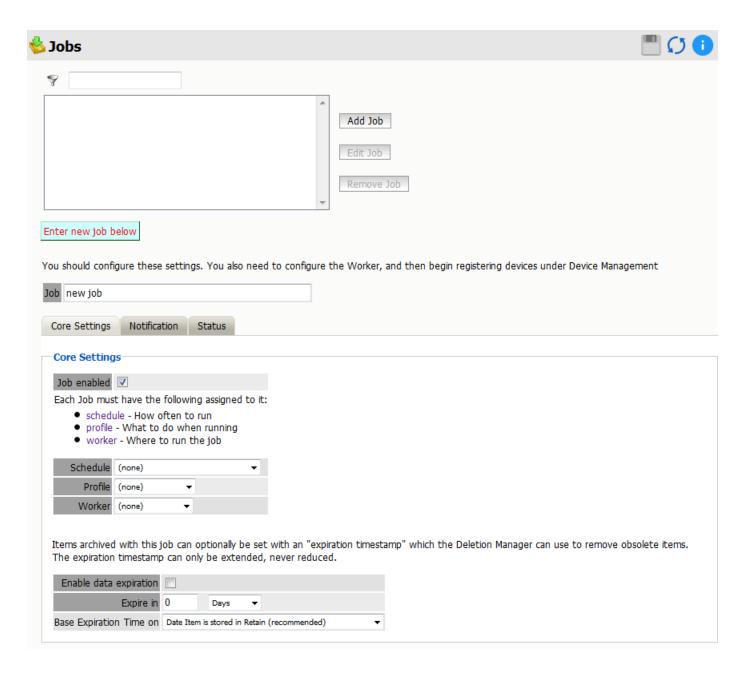


Devices which are active are listed along with daily statistics. The daily statistics are reset every night.

This page is mainly an informational page.

Core Settings

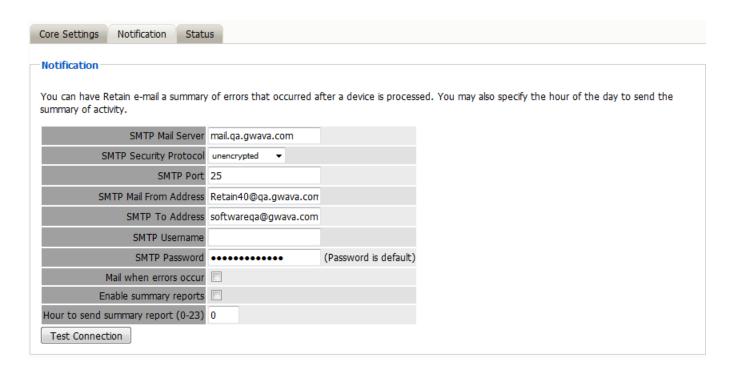
The core settings tab contains information on the schedule, profile, and worker utilized for the mobile module, as well as the data expiration date.



The Data Expiration is an option to place a time stamp on data in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, Caringo, nd Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one, incluidng Retain, may delete the item before its expiration date. Job Expiration is not retroactive for messages in the database, and only applies to messages archived by the job that it is active for. In order to have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

Notification

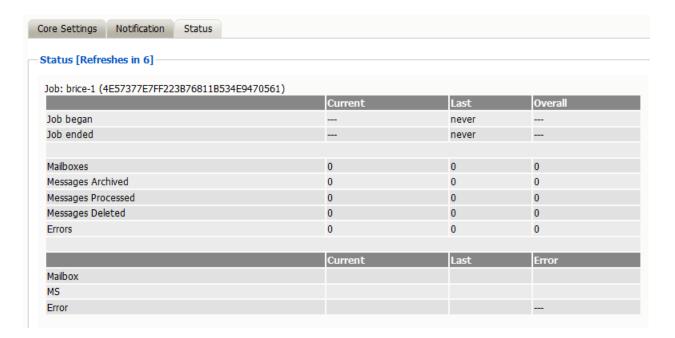
When a job is run, the notification option allows the administrator to be emailed a summary of each running job if desired.



For notification to function correctly, the SMTP information for the desired SMTP server must be fully filled-out. How much information is required varied depending on the mail system used.

Status

The Status tab displays the status of any currently running jobs, as well as the stats of the last completed job.



This tab is informational only for the Mobile module.

Mobile Registration Email variables

The Retain registration email template,

(...\RetainServer\WEB-INF\classes\config\mobile\registeruser.html), has many variables which allow for a customization of the registration email. There are a few variables which must remain unchanged, but the rest may be removed or modified, as desired, to adapt the email to any organization.

This is a quick description of the different variables, and which can be modified. They are distinguished by double brackets "[["and "]]". The Retain Server code that sends out the email sets these variables with the product specific and user specific information.

Typically, only #1 will want to be changed – for use in mobile device management software. If an admin wishes to make the registration email to look like it is coming from their company, they should modify #'s 1-2, 6-9. However, variables 6-9 may be simply removed if the admin doesn't want users to know what they are using.

- 1. Your install location This is typically the main variable to change. This variable is not included by default in the email to allow admins to utilize mobile device management software to push the Retain app to devices. It may be included in the email and filled in with the app's location so it can be manually downloaded to devices.
- Product This is automatically filled out with the Retain product name. Admins may set
 this to be the company name. This variable is used in the subject and the body of the
 email message.
- 3. Retainheader This is automatically filled with the Retain product logo. Admins may set this to any branding or company image file. This must be an image file.
- 4. Name Each user's individual username will be input here. This should not be changed.
- 5. GwavaURL Automatically filled-in with a link to the GWAVA website.
- 6. GwavaURLLabel This determines text displayed for the gwavaURL link.
- 7. ProductNameURL This is automatically filled-in with a link to the product website.
- 8. ProductNameLabel This determines text displayed for the productNameURL link.

The registration code and information in that variable must not change. This information is generated by the device management page when the device is added.

The following variables must not change.

Do not modify the follwing:

- enrollmentCode: This is generated when the device is created on the device management page.
- tenantid: This is static number for retain. Displayed on the mobile page from Module Configuration.
- email: This is from the Device Management page.
- name: This is from the Device Management page.
- endpoint: This is the address of the Retain Server
- rouerURL: This is pulled from the mobile page of Module Configuration.
- activateHREF: This is an identification string, a link specific to the user and installation.

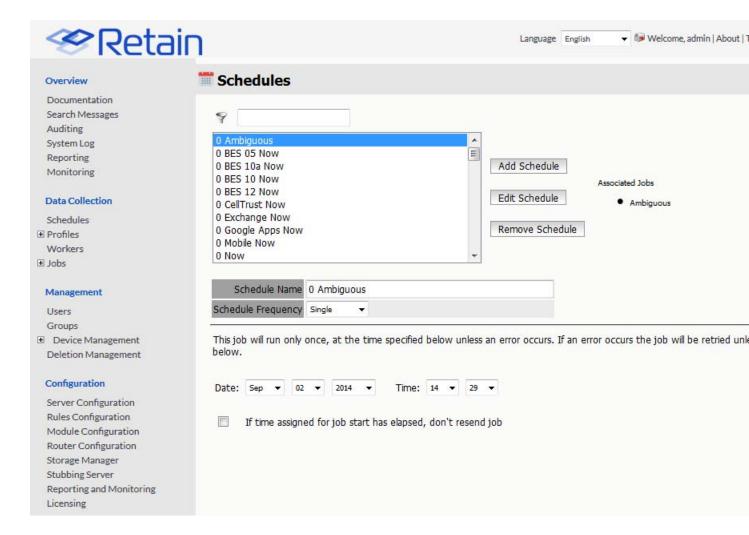
Schedules

Schedules define when to perform the data collection as defined in a job. a job defines what schedule, profile and worker to use.

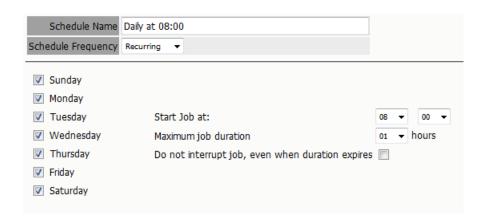
- You can create as many schedules as you need.
- Schedules have no functionality until they are associated with one or more jobs.
- They can be shared across multiple jobs.
- You cannot delete a schedule that is currently associated with a job.
- Schedules can be recurring or one-time.

Single schedules are configured with one start time, and one date.

To keep jobs from queuing up to the worker, select the box to do not resend a job that has an elapsed start time. Otherwise, a job will queue up on the worker if a previous job has not yet completed, causing the new job to start as soon as the current job is complete or terminated.



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



Choose when on which days you want this schedule to trigger a job.

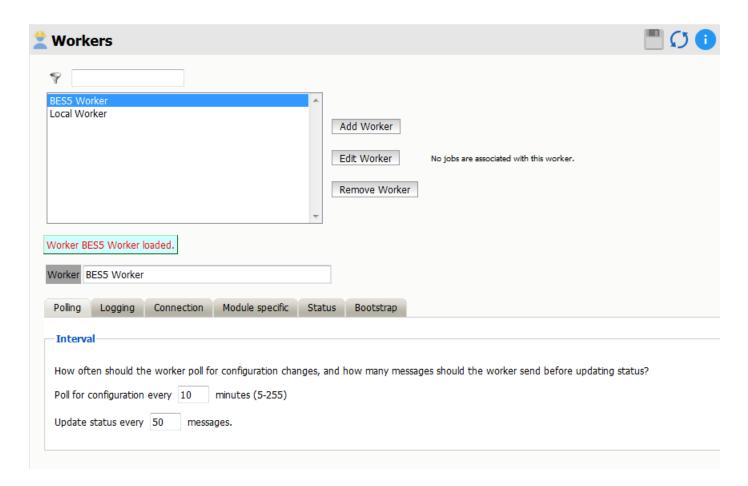
You may specify maximum job duration. If the job runs longer than the time limit you set, it will be interrupted. For cases where you do not want the job to be interrupted. In an example with a job you know will last longer than 22 hours, you may prevent the schedule from interrupting the job.

Don't forget to "Save Changes".

Workers

You will need to set up at least one Retain Worker. The Retain Worker does the actual work of collecting the data from a message system and delivering it to the Retain Server. The Retain Worker component is normally installed local to each GroupWise mail server, or any Exchange server from which you wish to collect data. Although a Worker may be installed on a machine running Retain Server, this is normally not recommended or sufficient, except for very small (100 or fewer users) systems or for evaluation purposes. Ideally, the worker you choose to dredge mail from a desired server is local to the server itself. Using a worker on Server A to pull data from a mailbox server on Server B is possible, but it will use up network bandwidth and will be slower than simply pulling data locally. For this reason, Retain was designed to allow you as many workers as you need, on Linux or Windows.

Worker Configuration Options

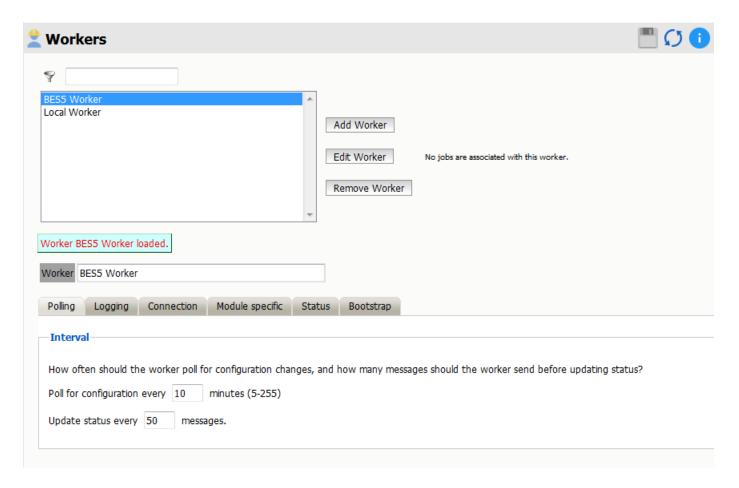


- Polling: Defines how often the worker polls for configuration changes. (Jobs, job changes, etc.)
 - o Default is 10 minutes.
 - o Default is to update the displayed status every 500 messages while running the job.
- Logging Controls logging by the Retain Worker.
 - o Default is normal.
- Connection. Important. This is the information the worker needs to connect to the Server.
 Double check this setting.
 - o Protocol

- o URL/host name
- o password
- o If the connection information is changed, you will need to manually copy the Retain Worker.cfg bootstrap file over to the Worker
- o Correcting or recopying the Bootstrap:
 - 1. Stop Tomcat.
 - 2. First delete the Retain Worker.cfg file on the Worker machine. For security reasons, there is no way to do this except by manually deleting the file from the server console itself. The path is: Retain Worker/WEB-INF/cfg. Removing the bootstrap resets the worker to initial configuration.
 - 3. Restart and re-upload the bootstrap to the worker.
- Status Basic status information on the Worker. This status information is updated while the worker is running according to the "Update Status" interval described above.
- Bootstrap Indicates the download link and when it was last downloaded.

Worker Configuration

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



- 1. Click "Add Worker".
- 2. Put the name of the worker in the "Worker box".
- 3. Example: Linux Worker 1.

4. Check the Connections Tab to be sure the Retain Server IP address settings are correct. This tells the Worker where the Retain Server is located. (Server Host Name or IP address must be correct; this should be populated automatically, but double check for accuracy.) The Server Connection setting must be able to resolve to the Retain Server if placed in a browser address bar.

Polling Logging	Connection Module	e specific Status	Bootstrap
Server Connection	n		
important to note t		e URL from the point o	highly recommended that you use a DNS host name for the Host Name. It's of view of the Worker contacting the Server , which may or may not match er.
	sing the Reload interface, etting of the Reload Serv	· ·	word assigned must be the same for each worker that contacts the Reload s password.
Server Protocol	http ▼		
Server Host Name	localhost		
Server Port	48080		
Server Path	/RetainServer		
Worker Password	•••••	(Password is default)	
Note: Changing and	of these parameters will	require recopying the	worker bootstrap file!
	, o. c. coo paramacara viii	. equil : coop; ing the	

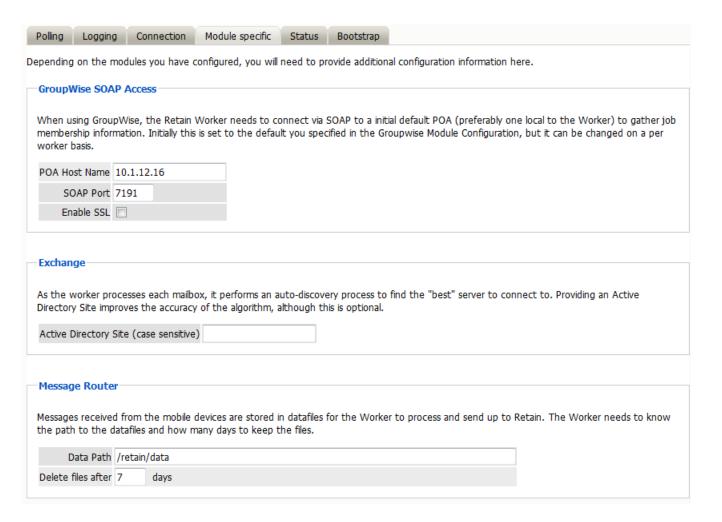
The Connection tab holds very important contact information for the worker's connection to the Retain Server and the Post Office. Make sure that the Retain Server connection address is correct, or the Retain Worker will not be able to talk to the server.

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

Important Note: It is vitally important that the connection information is correct before you save the worker. On a single server system, it is easy enough – it's the same server. However, in an enterprise setting where you might be using different components on different servers, it is extremely important that each component knows how to connect to the Retain Server.

If you do not set this correctly, the Retain Worker will not be able to contact the Retain Server and it will not run jobs.

- 5. The Module Specific tab contains information specific to individually configured modules. With Exchange, you may optionally specify a specific Active Directory Site the worker is located in. That allows Retain to find the "closest" Exchange Server that is appropriate for the user to be archived. The settings specified here are tied only to the specific worker configured, and does not affect the rest of the system. In general this setting may be ignored unless required. If required, input the appropriate information.
- 6. Message Router. The message router settings must be set on the message router worker to reflect where the data path is for the Message Router. This path must be specified from the worker's point of view. If the worker is not local to the router, then the path must be fully accessible to the worker; direct file access is required. You may specify how long to keep files in this location. Default is shown.

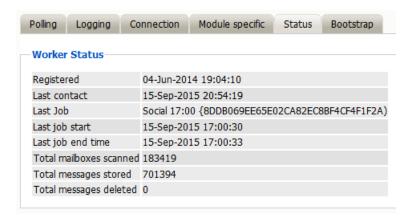


7. After you have checked and set your settings, click "Save Changes" before browsing to another page to finish creating and to save your worker.

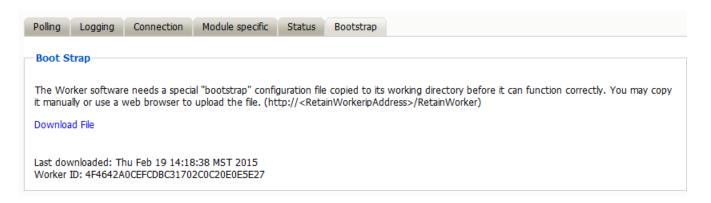
After the worker has been saved, two additional tabs become visible for settings on the worker. You MUST save changes on a newly created worker before the 'Bootstrap' tab is available to be used.

The status tab for the new worker will not hold any pertinent information until a job has been run or the worker has registered to the server and is in communication. Once a job has been completed, information will be displayed on the general status of the selected Worker.

Check this page after the worker has completed connecting to the Server for information and jobs. After a job has been run, or is running, more pertinent information, such as total mailboxes scanned and messages stored, will be displayed.



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



- 9. Save this file to your hard drive in a place where you can find it. Click "Save".
- Open a new browser window and type: http://< your-RetainWorker-ip >/RetainWorker
 The upload screen appears when a worker has not been initialized.
- 11. Upload the bootstrap file.



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.

Browse No file selected.	pload Bootstr	ap Configuration		
000 (394) 8276/3071354	Browse	No file selected.		

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.
- 12. Click "Browse" to find the file you just downloaded, and select it.
- 13. Click "Open".
- 14. Click "Upload File".
- 15. If the upload went well, you will see this screen:

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

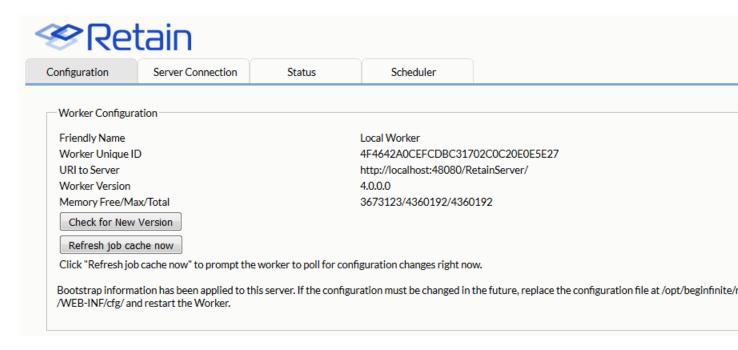
Continue

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

16. Login as Admin.

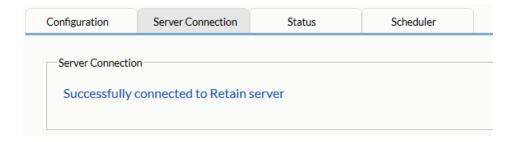


On all subsequent logins, you should be shown the configuration page first.



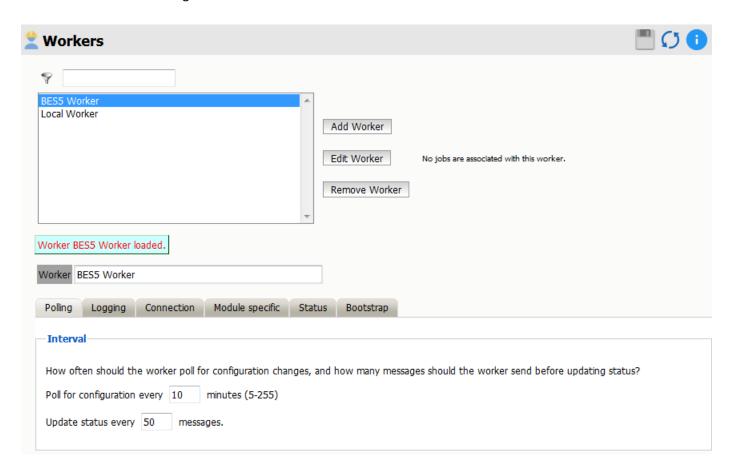
You should be able to see your configuration on the screen, as shown. The "Check for New Version" button checks for updates. The "Refresh job cache now" button tells the worker to poll the server for configuration changes and new jobs. After this button is selected, the Worker page will refresh automatically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

17. Verify the Retain Server connection success by selecting the Server Connection tab.



- 18. Go back to the Retain Server.
- 19. Click "Workers".
- 20. Click on the worker you just created.

You should see the following:



Pay close attention to "Poll for configuration every 'n' minutes. This is how often Retain will check for new jobs.

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

Now, set logging to a level you would like. "Diagnostic" writes extensive amounts of data to the logs. If using this level, only use it until satisfied that the system is fully operational. Please be aware, however, that using this setting will write a huge amount of data in the log files.

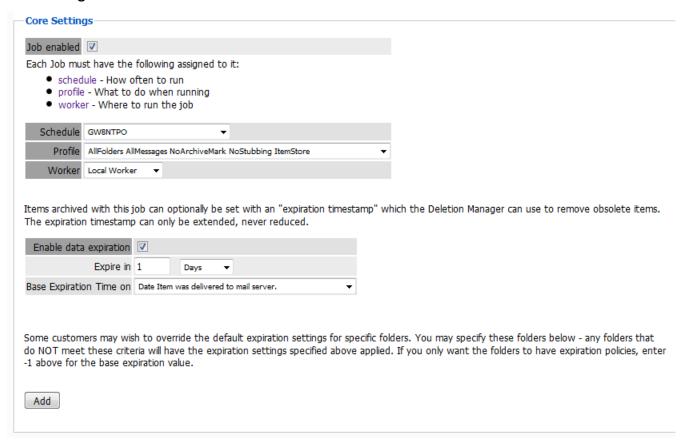
The profiles, schedules, and jobs must be setup, if they already have not been, in order to begin actual data collection.

Jobs

The Jobs section is where the Worker, Profile, and Schedule are all brought together and associated into a comprehensive configuration by allowing the user to tie their configuration to a selected SecureLine server and selected user list. The desired SecureLine server must be selected. The Worker will not start archiving anything until a created job is bound to a Worker, Schedule, and Profile.

Jobs, like profiles, will not be displayed until the module has been configured. Jobs are module specific. While the majority of the job interface is general to all modules, there are some specific settings.

Core settings

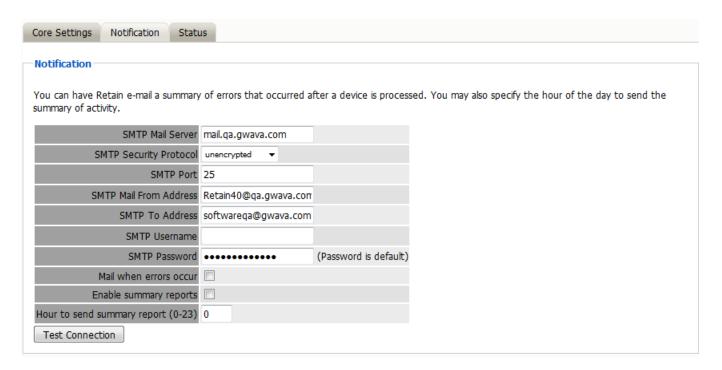


The Core Settings of a job contains configuration which must be set for the job to be saved and become active. A job must be enabled before it will run. Jobs must also have a specified schedule, profile, and worker. These are all selected from drop-down menus, and will not be populated unless those items are already configured in the system.

The Data Expiration setting is an option to place a time stamp on data in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, Caringo, and Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one (including Retain) may delete the item before its expiration date. Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for. In order to have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

Notification

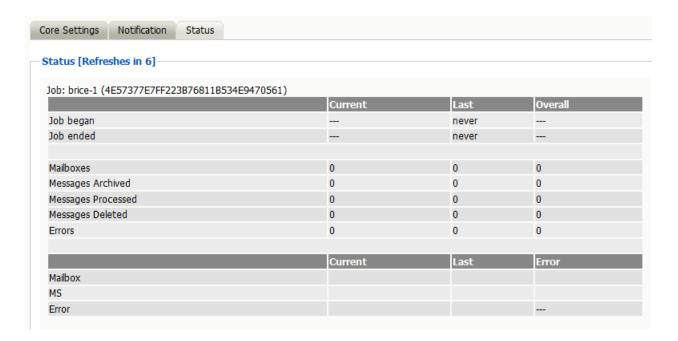
When a job is run, the Notification option allows the administrator to be emailed a summary and report of any errors, for each running job.



For notification to function correctly, the SMTP information for the desired SMTP server must be fully filled-out. How much information is required varies depending on the mail system used.

Status

The Status tab displays the status of any currently running jobs, as well as the stats of the last completed job.

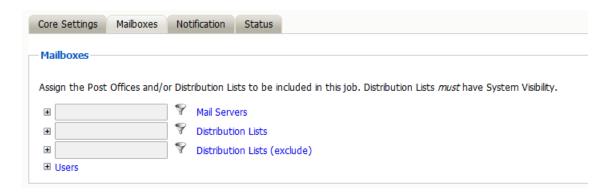


On some modules, currently running jobs may be terminated here. For the rest, this tab is informational only.

Mailboxes

The mailboxes tab is where the administrator specifies which entities (mail server(s) and/or Distribution List(s)) are to be scanned. This tab is not displayed for the mobile module.

Expand the Post Office and/or Distribution List trees, and check off the items you want to be dredged. Note: If you desire to have a job backup a single user, or selected group of users, select the Users menu and assign the user(s) desired



The Distribution List selection allows you to include or exclude a group of users from an archive job. If you want to use GroupWise Distribution Lists, the visibility needs to be set to "system wide".

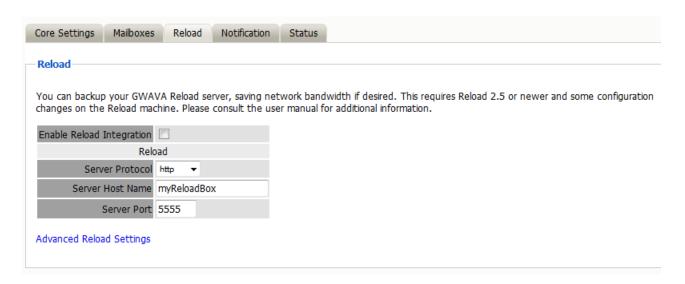
The users section allows you to select individual users to include, or exclude them from an archive job. For example: you can select an entire Mail Server to be archived, and then expand the users section to include or exclude users to the job.

This can also be used to select only certain users in the system for an archive job.

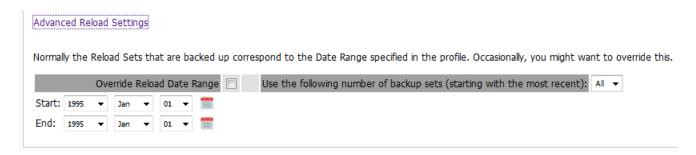
To add a user to the Include or Exclude list, select the respective 'Add user' button and search for the user. It can be helpful to unselect the 'only show recently cached items' option. Add the selected users to the list in the search window, then select 'Ok' to add them to the include or exclude list.

GroupWise Job option - Reload

Selecting this option tells the job to use the Reload integration for systems utilizing Reload. See the Reload Integration section for instructions on configuration. The Server URL must be edited with the correct contact info for the reload server. The default Reload port is 5555.

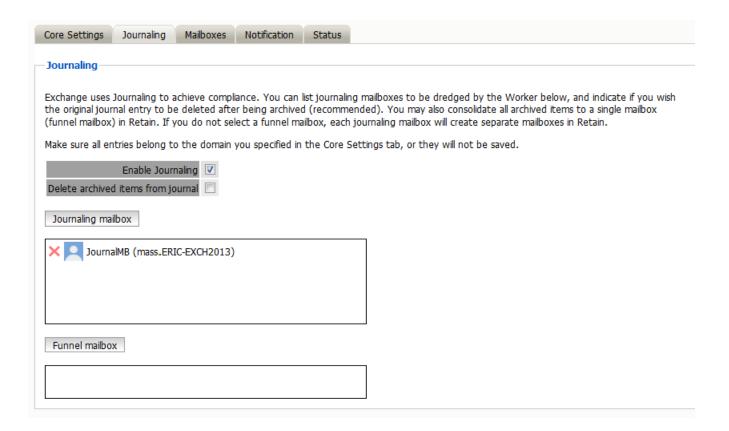


Retain instructs Reload to load the corresponding backup set for each date range in the profile. If a different set needs to be specified, the override setting specifies a date range specifically for Reload, which then will ignore the date range in the profile. Retain will still archive according to the profile date range or settings.



Exchange Job Option - Journaling

In order to achieve compliance, Exchange utilizes a Journaling mailbox. This mailbox can be set to be archived by retain to collect all messages on the system. The Journaling mailbox can rapidly grow in size if it is not cleaned out after messages have been archived. The Journaling option for Exchange jobs allows Administrators to set whether Retain will automatically clean out messages from the Journaling mailbox which have been archived.



On larger systems where there are multiple journaling mailboxes, Retain will automatically create a mailbox for each of them in the archive. However, if desired, the journaling mailboxes may be all archived to the same specified mailbox in the archive. This is the funnel mailbox. If desired, specify the mailbox by selecting the 'funnel mailbox' button, search for and select the desired mailbox. Only existing mailboxes in the Retain system may be specified as a funnel mailbox.

Management

Users

User and Rights Management in Retain include:

- Creating, deleting, and editing users
- Allowing new user accounts, and restricting specific ones from being created
- User expiration
- Assigning users to groups, to conveniently grant rights or set initial settings on a multiple user basis
- Granting access to mailboxes other than the user's personal mailbox
- Changing the specific functions the user can perform

To access User and Groups Management, the user logging in needs the "Manage users and groups" administrative right.

Creating Users

The primary purpose of a user account is to store their preferences, rights, mailboxes to which they have access, and authentication information.

Users come from one of two places

- They may be valid message system users logging in with their credentials
 - o These users use SOAP authentication for GroupWise and Active Directory authentication for Exchange; Retain checks their login credentials with GroupWise or Exchange
 - o These users initially belong to the group default. You may change this later
 - o You may restrict users (prevent them from logging in)
- Users may be specially created in Retain independently of any message system
 - o Users created in Retain do not need to have a message system account
 - o Users who don't exist in the message system will use the offline password

Offline Password

There are, however, occasions when you might want someone to search through the Retain archives, but who is not part of the mail system. Such a person might be an independent auditor, a lawyer, a user deleted from the live system, etc.

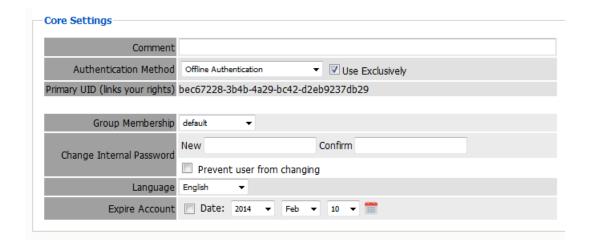
For this reason, Retain has an *offline password* system. These passwords are stored in Retain's control database. Retain does not care how a user authenticates: whether offline, via SOAP for GroupWise, Exchange, or LDAP, the same rights can be assigned. An administrator who possesses the Manage Users and Groups' administrative right can assign all the rights they contain.

Users may also be assigned access to more than one mailbox. Offline users will need to be given access to at least one mailbox to perform searches. Users who are assigned "Search All Mailboxes" rights have access to all users' mailboxes.

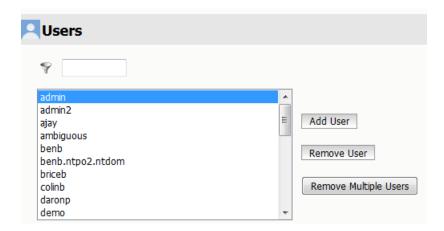
GroupWise Proxy support only works for users who authenticate via GroupWise SOAP protocol.

Creating a new user

1. Begin by clicking on the "Add User" button



- 2. Enter a new user name and then fill out the options under each tab
- 3. When you are done, click the save changes disk icon at the upper right



All previously created users are listed and can be edited or removed, by clicking on the user and editing the values, or through clicking on the Remove User button.

Core Settings

Authentication method

- SOAP (for GroupWise users)
 - o GroupWise users logging in are authenticated using SOAP
 - o These users are automatically entered into Retain's user list
- Exchange (SOAP for Exchange). Exchange users are authenticated via SOAP for Exchange and users are added into Retain's user list
- Offline Password (credentials stored within Retain)(any type of user)
- Google IMAP. Google users are authenticated through IMAP to the Google system.
- LDAP authentication The LDAP Authentication must be setup under Server Configuration | accounts tab
- You may lock a user account so it can only use one type of authentication
 - o If "exclusive" is not checked, it will try one, then the other

Primary UID

- Offline only accounts usually won't have one
- The initial admin account is set to use offline exclusively, so it never has one

Group Membership

- Default is "default"
- Create groups under "Groups" and they will appear as choices here
- Users may belong to one and only one group. From an assigned group, additional rights, mail-boxes, and initial settings may be inherited

Account Expiration

- Check this box if this user's account should not expire. Useful for administrator accounts
- By default, accounts never expire, (0=never), but can be changed in Server Configuration

Offline Password

- If you use this authentication method, store the password here
- May be changed as needed
- You can prevent the user from changing it themselves
- Passwords are always stored in an encrypted format never in clear text

Default Language

• Choose which language will be used in the Search Interface for this user

Disable account

• This allows the admin to pick a date when the account will no longer be allowed to login, but the account will not be deleted

User Rights

The restore r

The restore right has been activated. Please understand how this functions for your messaging system(s).

When restoring to Exchange, the full item is restored, including attachments. There are no dependencies.

In GroupWise, we restore just the "stub" of the item. This has the advantage of being quite quick, and accurate.

However, for your users to read the stubbed items, you MUST enable the Retain Stubbing Server, and must set up stubbing in ConsoleOne. You do NOT have to enable stubbing while archiving messages.

Rights explicitly granted to the user.

Administrator level rights

- Access Reporting and Monitoring Server
- Access all audit logs
- Deletion Manager
- Add, edit, remove global tag definitions
- Apply or remove litigation hold
- Manage Server
- Manage Users and Groups
- Manage Workers, Schedules, Profiles, Jobs
- Publish Messages
- Restore messages [Any Mailbox]
- Search all mailboxes
- See Confidential Items [other mailboxes]

<u>User level rights</u>

- Apply Confidential Tag [other mailboxes]
- ✓ View/Save attachments
- ✓ View personal audit log
- Delete messages in other mailboxes
- Delete personal Messages
- Export Messages
- Forward Messages
- Print Messages
- Read Configuration (Redline)
- Restore messages [My Mailbox]
- Apply Confidential Tag [own mailbox]
- Add, edit, remove user tag definitions

The following rights are additionally inherited from the user's group membership:

View/Save attachments Forward Messages Print Messages

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
 - o You don't need any of them for the user to access their mailboxes
 - o 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

Administrator level rights Access Reporting and Monitoring Server Access all audit logs Deletion Manager Add, edit, remove global tag definitions Apply or remove litigation hold Manage Server Manage Users and Groups Manage Workers, Schedules, Profiles, Jobs ✓ Publish Messages Restore messages [Any Mailbox] Search all mailboxes See Confidential Items [other mailboxes] User level rights Apply Confidential Tag [other mailboxes] ✓ View/Save attachments ✓ View personal audit log Delete messages in other mailboxes Delete personal Messages Export Messages Forward Messages Print Messages Read Configuration (Redline) Restore messages [My Mailbox] Apply Confidential Tag [own mailbox] Add, edit, remove user tag definitions

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 Rights Summary

Administrator-level rights

Access Audit Logs

- The Audit Logs become visible
 - o Main menu audit log
 - o Message properties audit log
- You can see user activity on two levels
 - o Globally
 - o On an individual message basis
- Deletion Manager
 - o Allows management of Deletion jobs
- Add, Edit, Remove global tag definitions
 - o Allows manipulation of global tags in the view messages interface
- Litigation hold
 - o Allows excluding mail from Deletion
- Manage Server
 - o Access Server Configuration and Licensing
 - o Access diagnostic utilities
- Manage Users and Groups
 - o Access Users and Groups
- Manage Workers, Profiles, Jobs, Schedules
 - o Access data collection system
 - o Login to Worker web interface
- Allows use of Publisher
- Allows messages to be restored to all mailboxes
- Search All Mailboxes
 - o Grants access to all mailboxes in the system
 - o Allows the user to search and browse them
- See Confidential Items [Other Mailboxes]
 - o Allows users to view items which others have tagged as confidential

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

User-level rights

All user level rights are strictly optional, and add additional functionality.

- None are needed to access your own mailbox and other mailboxes assigned to you
- You may wish to grant Forwarding, View Attachment, and Printing rights
- The initial "default" group created upon installation does this

Note:There is no way to perfectly block printing in a web browser, so using this feature should not be taken as a 100% guarantee that users won't be able to print. Nonetheless, for most users, it is effective

These rights are self-explanatory:

- View/Save Attachments
- View own audit log
- Forward Messages
- Restore Messages (My mailbox only)

Print Messages

Other rights:

- Delete messages in other mailboxes
 - o Allows access to delete items from mailboxes that the user is granted access to, but is not their personal account
- Delete personal Messages
 - o Allows deletion of items from the user's personal account
- Export Messages
 - o Grants the right to use the export option which sends groups of messages or items as a single PDF document.
- Add, Edit, remove user tag definitions
 - o This allows the user to create, remove, and edit the user's tags in the search messages interface.

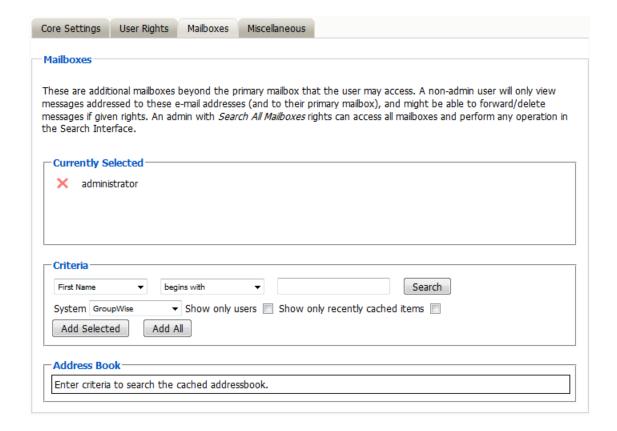
The Read Configuration right (Redline Integration)

If you are integrating with GWAVA's **Redline** monitoring product, you will need to create a user account so that Redline can log in and retrieve monitoring information. We recommend the following settings:

- Account Never Expires
- Offline Password Authentication is required. (use exclusively) (be sure to set the password)
- · Read Configuration (Redline) right.

Mailboxes

Select the mailboxes this user will be able to access in addition to their own.



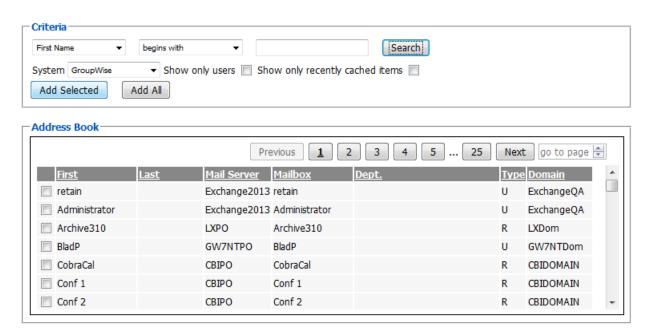
Often, you will want some users to be able to search through more than just their own mailbox. Administrators can have "Search All Mailboxes" as a right which gives them access to everything. This section allows you to give a far more selective range of mailboxes to a user for searching.

In the example above, the user has explicit rights to two mailboxes. These mailboxes can be taken away from the user simply by clicking on the red 'X'.

Adding users to the list is done using the Address Book selector. In the criteria section, you may enter information to search for a mailbox or a set of mailboxes. The search results will appear in the Address Book section. Each listed entry has a check box you can use to select that mailbox for addition to the list. Once you are done selecting, click Add Selected Items to add those mailboxes to your list of searchable mailboxes.

Address Book Selector

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



It shows the currently selected items at the top, and lets you delete an item by clicking the red X. (The New Mailbox selector in the Search Interface is an exception; just choose another item)

Adding Mailboxes

- 1. Choose between the configured module systems
- 2. Fill out basic criteria to narrow your search results (or no criteria for the first 100)
- 3. Click Search
- 4. The results up to a maximum of 100 are displayed
- 5. The user can then page back and forth among the first 5 pages of results
- 6. Choose which of the results you want to add to the selected list
- 7. Click Add Selected Items

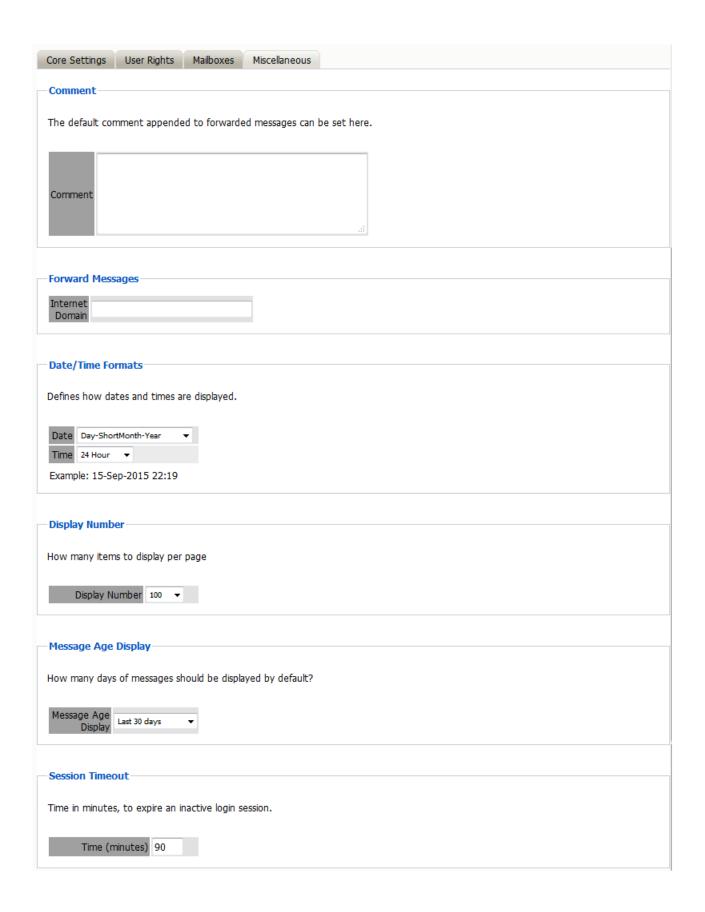
Notes: You can restrict to just Users (skipping Resources). You can show only recently cached items (last 10 days). The search is not case sensitive.

About "Show only recently cached items"

This option restricts the list of items shown in the selector to those with items stored within the last 10 days. In user/group management, it restricts the list to users who have logged in to the live Mail system within the last 10 days. The idea is to show only current items. If you DO want to see all items regardless of whether they've shown activity within the last 10 days, just uncheck this option.

Miscellaneous Tab

This tab contains settings that mainly govern the way the Search Interface works for the selected user. Note that the user can change any of these settings by using the User Options tab in the Search interface.



- Comment Default comment for forwarding messages
- Date/Time FormatHow you want your dates and times to be displayed
- Display Number How many items to display per page
- Forward MessagesAutomatically append the specified address to forwarded messages
- Message Age DisplayDefault date filter for searching. May be changed on the fly
- Session TimeoutInactive session timeout. Can be between 5 and 60 minutes

GroupWise Proxy Support

Retain supports the GroupWise proxy function. To enable it, check the box in the Module Configuration section. (NOTE: using proxy is useless if the user you wish to enable this function for is set to use offline authentication – found under the core settings of the user)

note: the 'all user rights access' in groupwise is not supported.

This function is used to enable a user to access the mailbox of another user. For example, if user B grants the right to user A to access their mailbox in the GroupWise client, then user A can "proxy" in to user B's mailbox.

Much the same way, if user A has proxy rights into user B's mailbox in GroupWise, and the function is enabled in Retain, then user A may select user B's mailbox for browsing or may search through user B's mailbox in the Search Screen.

In Retain, it is the MAIL READ right which grants access.

Retain uses the list of available mailboxes shown in the GroupWise client to determine which mailboxes will be made available to the logged in user (user A in our example). Thus, it is important that user A has logged into user B's mailbox as proxy using the GroupWise client before doing this in Retain. While user B might have granted the rights to user A, if user A has not yet logged in as proxy to user B's mailbox with GroupWise, then user B will not appear in user A's list of available accounts to proxy into.

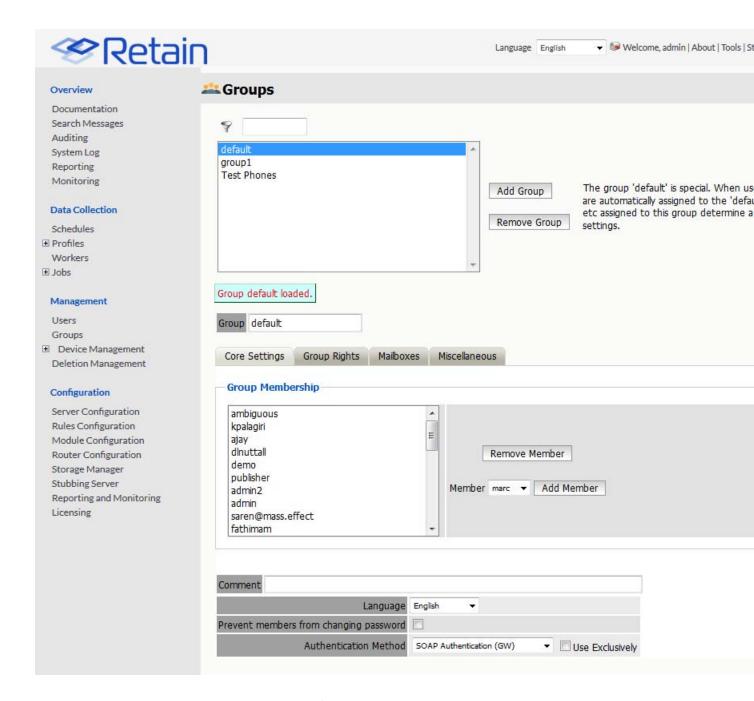
Retain checks these proxy rights the first time you access a proxy users mailbox, then caches the information for 7 days as configured in the server Configuration – Miscellaneous tab. (Default is 7 days.)

If you have access to another mailbox by virtue of GroupWise proxy, then you will see that mailbox appear in the mailbox selector in the search screen or you may search through that mailbox as well.

Groups

Once you have created users, you can group them and give them common rights. For instance, you can make a group called "Auditors" and this group will have rights to certain specific functions. If you then add users to this group, these users will get those specific rights in addition to the rights you have explicitly given them.

This is a great way to save time in managing the rights of a large number of users.



This is the main Groups administration screen. It's laid out just like the Users screen and rights are assigned the same way, for the most part.

In the Groups screen, you will see all defined groups .The group *default* is created when you install Retain. Users are automatically assigned to this group. If you want users to be part of another group, you must first create that group and then add users to it.

You will notice the drop-down list that allows you to select a member. If it says "(no users)", it means that all users are part of the currently selected group and that there are no more users available to add.

You may specify the default language used in the Search Interface for this group of users.

Creating a New Group

If you want to create a new group, simply

- 1. Click "Add Group"
- 2. Type the name in "Group"
- 3. Click "Save Changes"
- 4. Change the properties of the group in the tabs below

Group Rights

Group Rights	
Rights granted to the group. Members of the g	roup will inherit these rights, as well as ones explicitly assigned to them.
Administrator level rights	
Access Reporting and Monitoring Server	
Access all audit logs	
Deletion Manager	
Add, edit, remove global tag definitions	
Apply or remove litigation hold	
☐ Manage Server	
Manage Users and Groups	
Manage Workers, Schedules, Profiles, Jobs	
Publish Messages	
Restore messages [Any Mailbox]	
Search all mailboxes	
See Confidential Items [other mailboxes]	
Uses level stables	
User level rights	
Apply Confidential Tag [other mailboxes]	
✓ View/Save attachments	
View personal audit log	
Delete messages in other mailboxes	
Delete personal Messages	
Export Messages	
 ✓ Forward Messages ✓ Print Messages 	
Read Configuration (Redline)	
_	
Restore messages [My Mailbox]	
Apply Confidential Tag [own mailbox]	
Add, edit, remove user tag definitions	

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.

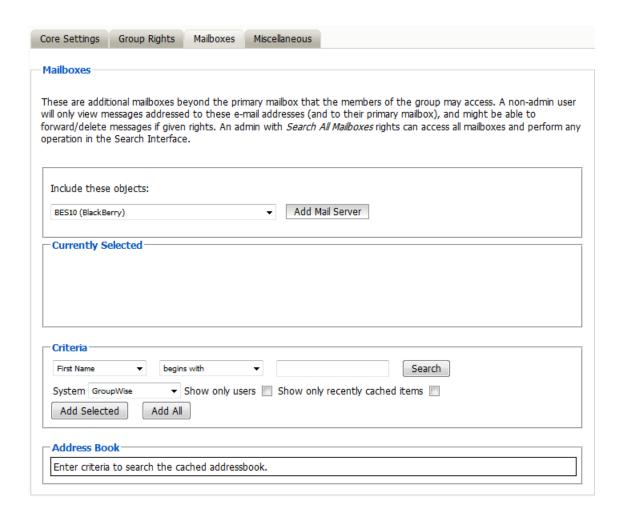
TIPS

Typically, you will not want to grant Redline rights to a group, but in a situation where you have more than one Redline control center and you want to see which one is monitoring, group-level Redline rights would be appropriate.

Typically, you will want to make the default group's rights rather restrictive, granting very limited access by default. If you wanted a user to have more rights, you would simply assign them to another group or you could assign additional explicit rights.

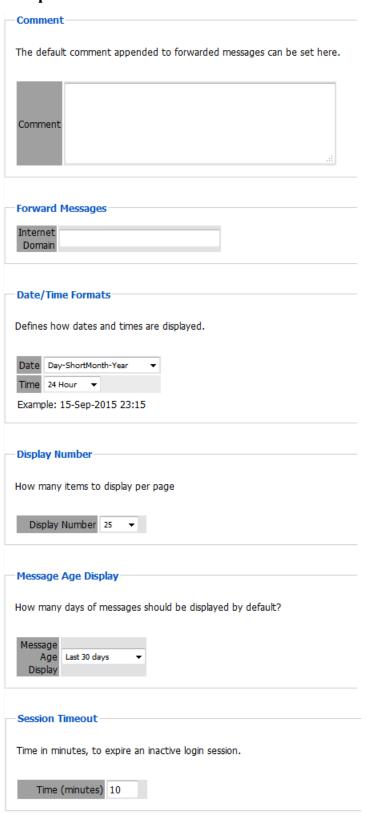
By using groups, you can have groups of individuals with rights to totally different sets of archives. Users only get rights to the mailboxes you allow them.

Mailboxes



- This screen works exactly like the user's access to mailboxes. Please see the user's mailbox section for details on how to select which mailbox or mailboxes to assign to the group.
- You use the address book selector to choose which mailbox or mailboxes to give the group access
- By default, groups have access to NO mailboxes; users may only access their own mailboxes. If
 the user exists only in Retain (no GroupWise account) and their group has access to no mailboxes, then they will have access to no mailbox either.
- Typically, you'll give group access to mailboxes only when you want all members of the group to
 be able to access a particular mailbox. Some examples could include a general sales account or
 accounts being audited by a group of auditors, a workgroup needing to access each other's
 archives, and so on.

Groups - Miscellaneous



This is exactly like the users' miscellaneous configuration screen. It applies to all users in the group and users' explicit settings will override the group settings.

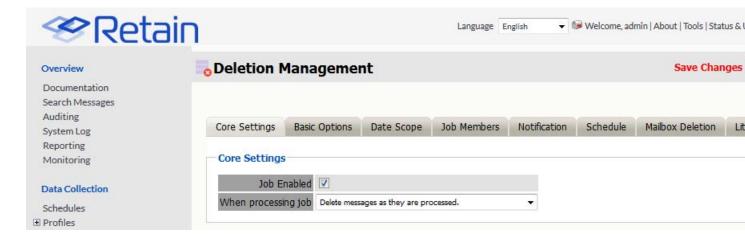
These settings govern the way the search interface works for the selected group. See more information about these settings in the users' miscellaneous section.

Deletion Management

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

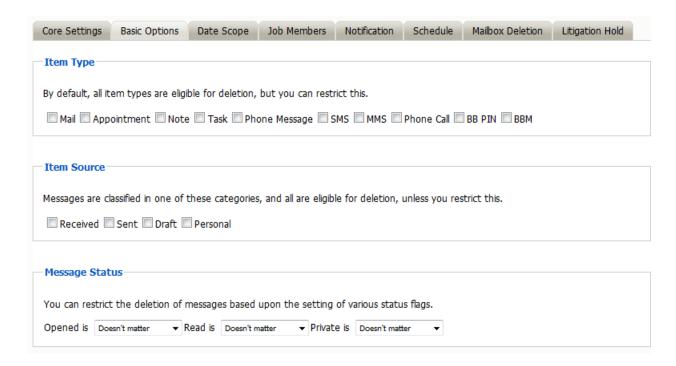
Core Settings

Here you enable and disable deletion jobs and detail which actions they will take. When setting up a deletion job, you have the option to tell the job to generate a report of messages that will be deleted, (the report is limited to 10000 items), or to delete messages as the job runs. The report function can be very handy to ensure that a job will not remove any message that is needed, before the messages are actually deleted.



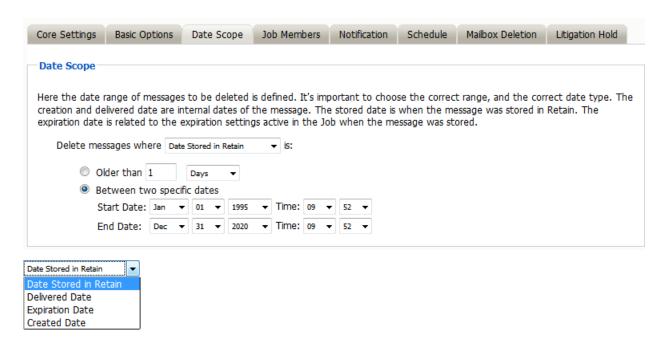
Basic Options

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. If all boxes are unchecked, the job will get everything.



Date Scope

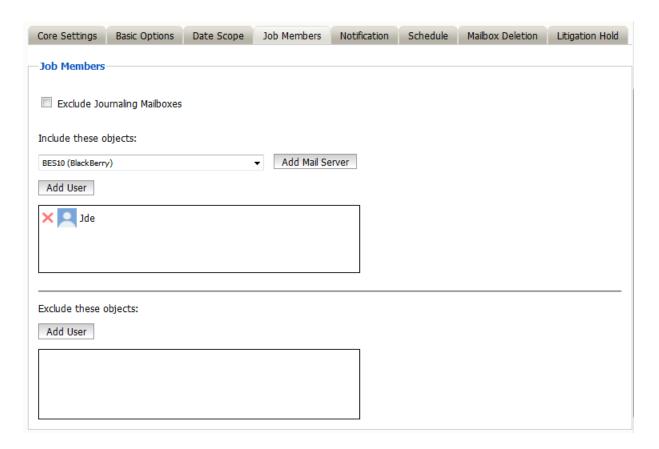
There are many dates that are contained in a mail system, and the deletion manager allows you to select different date ranges to identify the scope of the deletion manager. The setup is simple; the date range between the "Begin" and "End" dates will be targeted by the deletion job.



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

In order to have data with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete data with an Expiration Date older than 1 day.

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

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

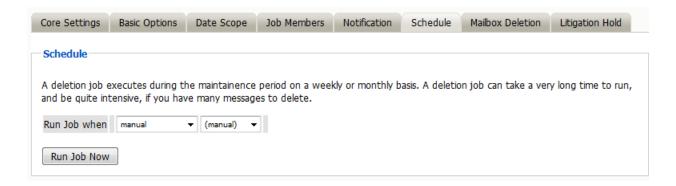
Notification



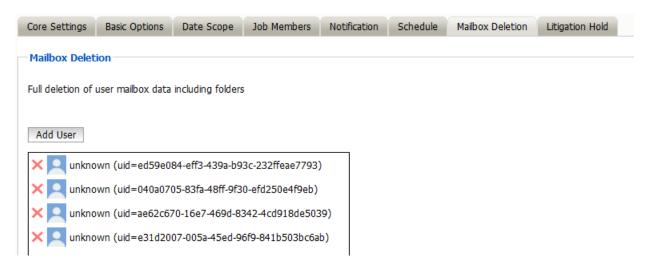
The reports, errors, and summaries of deletion jobs can be sent to the listed address in the notification tab. Using the Generate Report option in Core Settings is useless unless you set this variable. Select the options as desired.

Schedule

The Schedule 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, on a specific day of the month, or only when manually started. 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.) To manually start a deletion job, simply press the 'Run Job Now' button. This will immediately initiate a deletion management job.



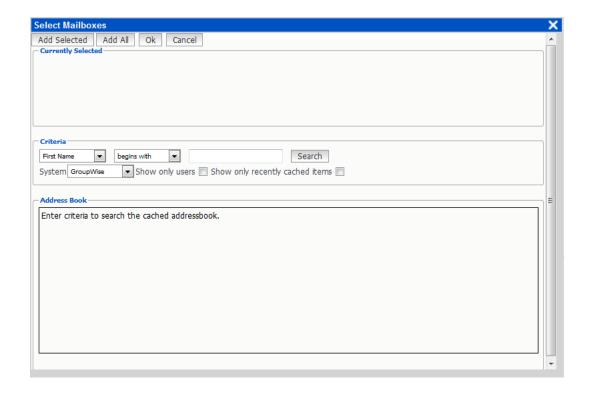
Mailbox Deletion



Deletion Management also has the ability to clean out mailboxes which no longer need to be archived. Mailbox Deletion will not remove the mailbox itself, but will remove everything in a mailbox; messages, items, and folder structures. Any mailbox which has a litigation hold applied to the whole mailbox will be skipped if placed on this list. Likewise, and any item or message with a litigation hold applied will remain with the dependent folder structure. Everything not protected by a litigation hold will be deleted from the mailbox, archive, and index.

Mailbox Deletion may be performed with a standard deletion job, but it is not something usually left with a scheduled, recurring job, because there is no need to re-delete an empty mailbox. It is recommended to use the Mailbox Deletion option on single instance jobs.

To set a mailbox for deletion, select the 'Add User' button to open the 'Select Mailboxes' dialog.



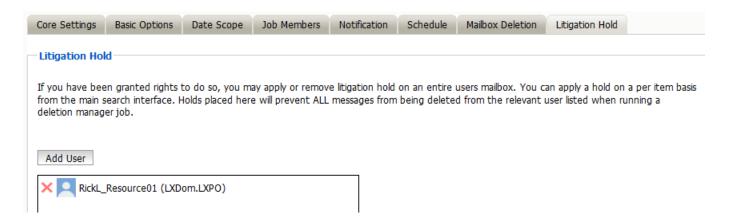
Search for and select the desired user, or users' mailboxes.

Select the 'Add Selected' button to add the selected users to the selected user list. Additional searches may be performed after selected users are added to the currently selected list. Repeat until all desired users are selected, then select 'OK' to add the user mailboxes to the deletion list.

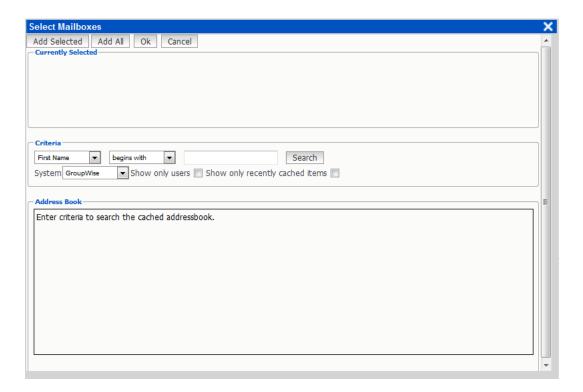
Save the changes on the page. There is no other work to be completed. Once the deletion job executes, the mailbox is cleaned out and all items not subject to a litigation hold are deleted.

Litigation Hold

The Litigation Hold tab provides the ability to exclude any specified user's data from any deletion job, preventing any of their data from being deleted when the job runs.



To add a user to the litigation hold list, select the 'Add User' button to open the 'Select Mailbox' window.

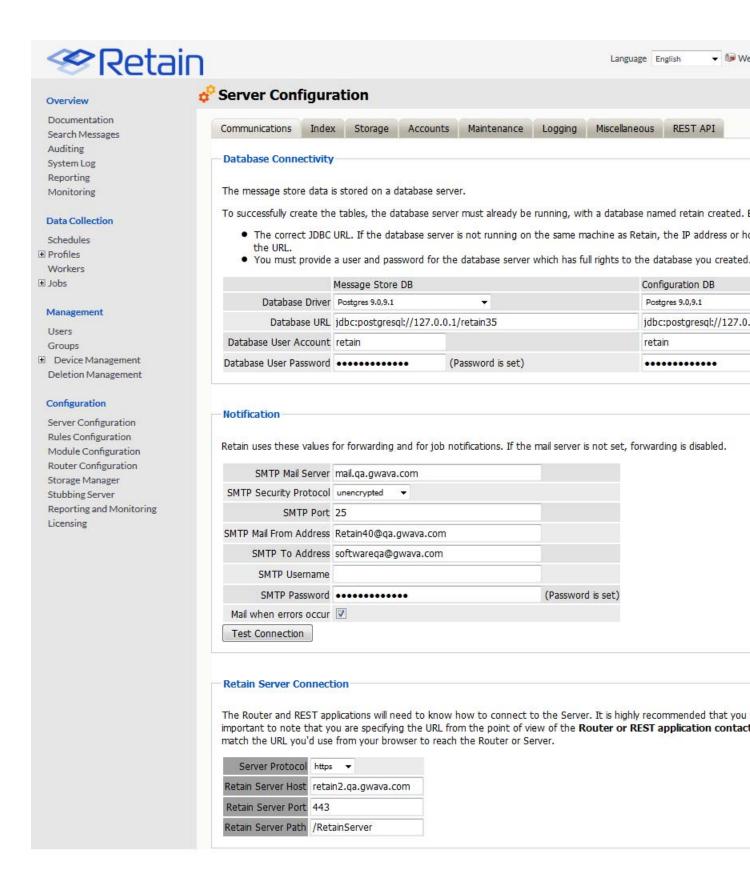


Select the source system for the user and enter search criteria. After searching, select the desired user or users and select the 'Ok' button to add them to the list. Save all changes.

Configuration

Server Configuration

The overall operation of Retain is configured in this section.



Along the top, you will see configuration options for the Retain server and workers known to this server.

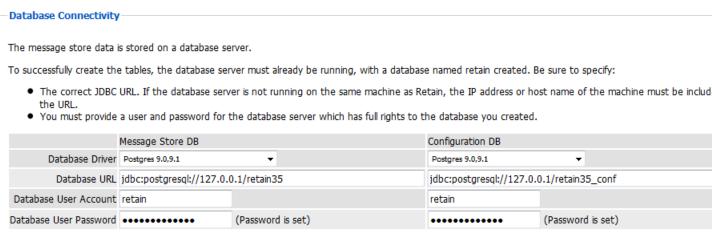
The most important topic is communications. These are the settings you have set when you initially installed Retain. You may make changes here.

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

Communications

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

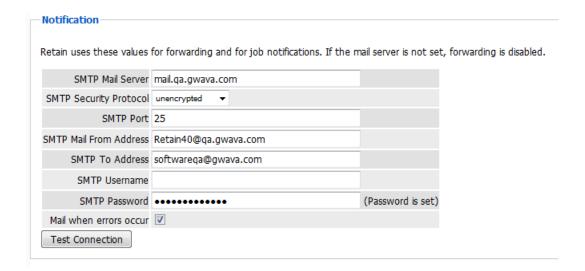
Database Connectivity



The Database Connectivity tab contains the connection information for both the configuration and storage databases. This information should never change unless the database server is being migrated to a new location.

Notification

In this location, the admin user specifies what address receives notifications about the general system, as well as what SMTP Gateway to utilize to send these notifications.

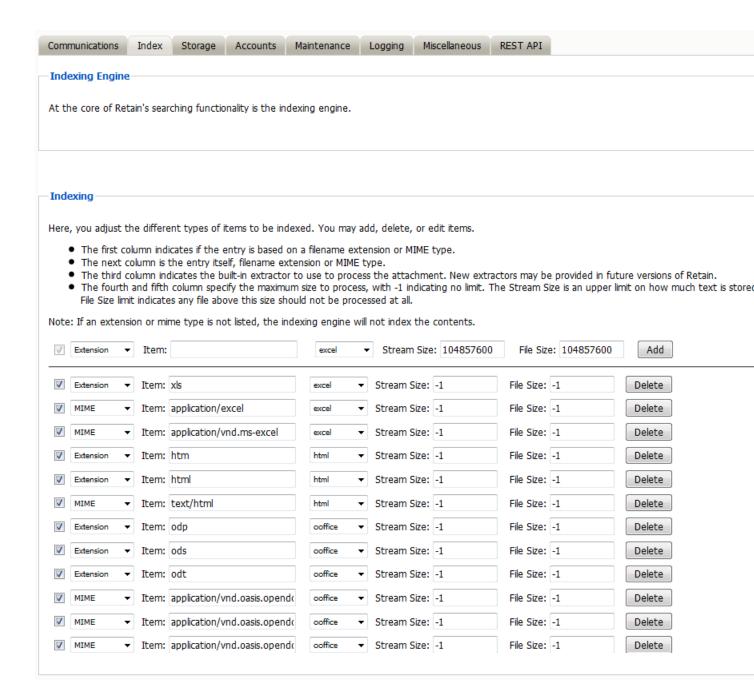


Retain Server Connection

This location dictates the connection address, port, and protocol to which the Retain Server will listen for communications from the Router and from any other REST applications. The communication settings must be correct from the point of view of the applications trying to use it. The hostname or IP address must successfully connect to the Retain Server.

Index

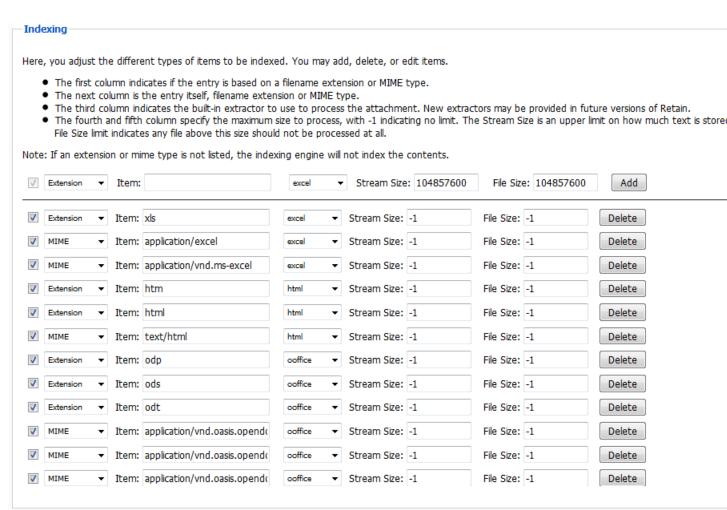
The Index tab allows the management of the indexing engine as well as the ability to decide what kinds of attachments are indexed and what size.



Retain utilizes two index engine choices. These choices are set at install or upgrade. The index engine is the same between the Standard and the High Performance Index options, (HPI). However, while the Standard engine is nested with the Retain Server, the HPI is a cluster of index engines, managers, and replicas which provide scalability and disaster insurance. The standard index engine will perform better than the HPI under all but heavy loads. (Index engine speed is limited by disk IO for the Standard, and Disk IO and network speed. The trade off of performance and bottleneck moved in the favor of HPI when the system is very large and under large demand.

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. Choose whether to index the attachment based on its filename extension or its MIME type (the content itself). You also choose which extractor to use to index the attachment. Select as many as you need. If an attachment type is common in the system and needs to be indexed, but is not already indexed, it may be added by using the 'add' row.



Storage



The default storage path for your Retain archives is listed here.

Retain uses a hybrid storage system. All message headers and metadata is stored in the database but all message text and attachments are stored as files on the file system in the location listed here (and subdirectories). This is done a little like GroupWise, where the OFFILES directory stores Binary Large Objects (anything larger than 2KB). In the case of Retain, this is transparent to the user and you may choose whether to encrypt this data. If it is stored in a location accessible ONLY to the Retain Server, then encryption is not strictly necessary.

Accounts

Communications	Index	Storage	Accounts	Maintenance	Logging	Miscellaneous	REST API			
— Account Manag	ement—									
, iccount i ianag	Cincine									
You may expire u	nused acc	ounts, disable	e new accoun	its from being (reated autom	atically, and restric	ct some users fr	rom loggi	ing in.	
Expire unused	accounts	after how m (0	any days =never)							
	[Disable new a	accounts							
		Prohibite	d Logins					A	Address	Remove selected addres
You may enable a password strength policy if you want to enforce strong passwords for users. Keep in mind these passwords may be overridden by individual messystems. Password strength: Low										
Must be between	n 5 and 15	characters i	n length.							
Intruder Lockout										
Enable this featur restarting Tomcat		orarily lockou	t an account	that may be th	e subject of a	n illegal entry atte	empt. You may	clear lock	kouts using	the button below, or b
	En	able Intrude	r Lockout? 🔽	•						
	Number o	of invalid logir	attempts 2	▼						
			(minutes) 5							
If triggered, lock Clear lockouts	account fo	or this period	(minutes) 5	•						

Open System vs. Closed System

Normally, Retain lets all mail system users log in. This is considered to be an "open" system. When that happens, Retain will check to see if a Retain account already exists and if not, it will create a new account for them and assign them to the group *default*.

Sometimes, you don't want certain users to have access to the Retain archives. In this case, you may add these users to the list of Prohibited Logins. You do so by entering their name in the Address field and click "Add".

To make a "closed" Retain system, simply click on "Disable New Accounts". If you use this option, it means that you will have to manually create accounts in Retain for authorized users. In other words, the only people who can access your system will be people who you specifically create an account for.

In Retain, user accounts expire after 30 days of inactivity by default. You may choose the number of days or choose 0 for "accounts never expire".

See "User Rights" for more information.

Password Strength

User-created passwords may be controlled for strength. By default, Retain accepts any password set by users. To require a higher security password, select the higher level desired. Requirements for the low, medium, and high settings are defined as:

Will accept any password

Will accept any password

Low

Must be between 5 and 15 characters in length.

Medium

Must be between 5 and 20 characters in length, with at least 1 lower case characters, at least 1 upper case characters and at least 1 numerical characters.

High

Must be between 8 and 20 characters in length, with at least 2 lower case characters, at least 2 upper case characters, at least 2 numerical characters, and at least pecial characters. Also, the password will be checked against a dictionary.

Intruder Lockout

Intruder Lockout			
Intruder Lockout			
Enable this feature to temporarily lockout an accou restarting Tomcat	nt t	hat	may be the subject of an illegal entry attempt. You may clear lockouts using the button below, or b
Enable Intruder Lockout?	V		
Number of invalid login attempts	2	•	
Time interval (minutes)	5	•	
If triggered, lock account for this period (minutes)	5	•	
Clear lockouts			

Accounts may be locked if multiple failed attempts are detected, according to the specified time window. This is useful to deny password cracking attempts on the server.

To enable Intruder Lockout, select the checkbox next to the 'Enable Intruder Lockout' option and save the changes. All changes will be immediate as soon as the save button is selected.

If a user has locked their account and requires immediate access to the system, all lockouts may be cleared. To clear any locked accounts, select the 'clear lock outs' button at the bottom of the page. There is no need to save changes; the clear command is immediate.

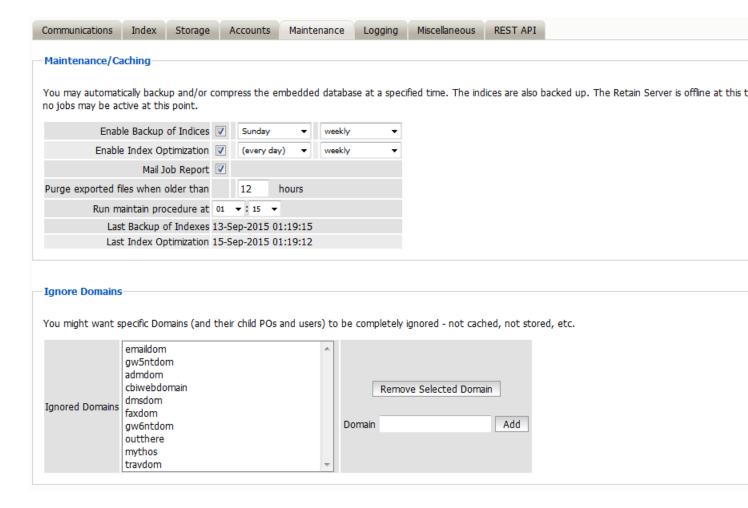
Maintenance

Maintaining a Retain system involves many tasks:

- Backup Configuration (Embedded) Database
- Compress Configuration Database

• Indexing (All Indexes)

These items are all controlled here.



This configuration database – the "Embedded" database:

- Should be backed up regularly
- Contains your system address book
- Also contains key information used by Retain
- Configure how often it gets backed up here.

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

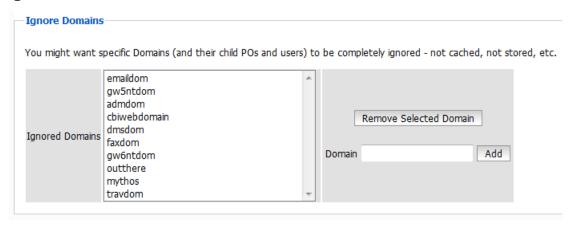
Retain also lists a history of backups and maintenance here.

Note: This maintenance section does not involve backup or maintenance of the data store. You must back that up yourself with your normal backup tools.

You must backup:

- SQL Database
- File System where your data is stored

Ignore Domains

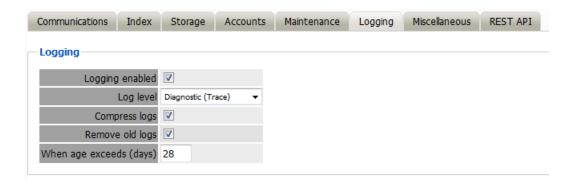


Some mail systems allow for the addition of non-system, or external domains. Sometimes, you might do this to add external domains to the mail system address book. However, the mail system cannot pull up e-mail from these domains and you may want to tell Retain to ignore references to these external domains to avoid wasting time during data collection.

Specify these domains in this window.

Logging

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



You also specify how long to keep logs here. By default, logs are deleted after 10 days. Logs may be 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.
_		, because they would generate enormous audit logs. Enable them with care. Also, consider clearing you he audit table on your SQL Server.
Days before audit records	are removed (0=never remo	ve) 0
Advanced auditing options	į.	
Select/Unselect all		
Changed User/Group	Changed Data Collection	✓ Searched
Mailbox Switch	Audited	Server Config Change
Exported	Published	Restored Message

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



You may choose the default worker password here. One is automatically generated for you when you create a worker but you may change the password here if you like.

Export XML

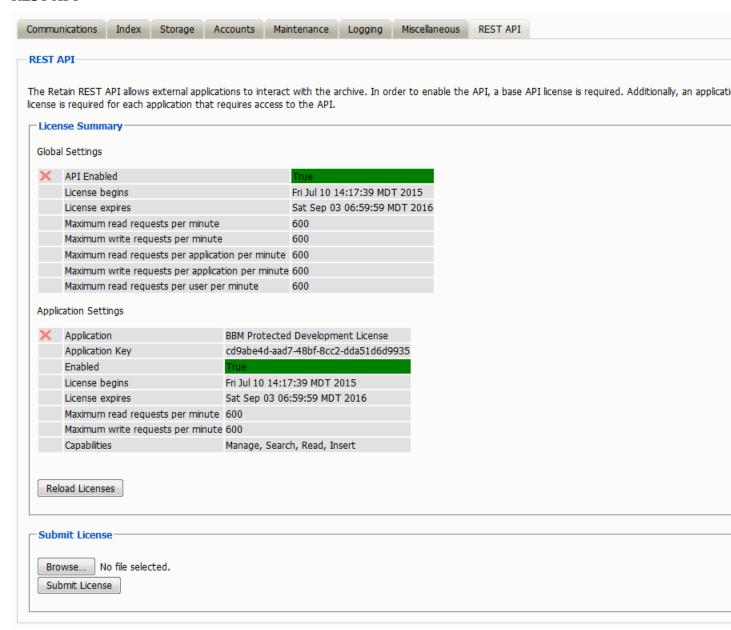
This XML Export function is included in Retain in case you have an XML compliance mandate. You enable it here.

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

Ordinarily, you would not want to do this because it consumes enormous quantities of disk space, loads up your file system and degrades performance.

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

REST API



The REST API is an interface built into Retain which allows third-party applications to input data into the Retain archive. This API was developed to open new horizons to the Retain Archive for any application which conforms to the API requirements. Because of the access the API grants to applications, an additional API license file is required for each application. This license contains credentials and access rights. In order for a third-party application to utilize the REST API, the application must have a credentialed key provided by GWAVA Inc.

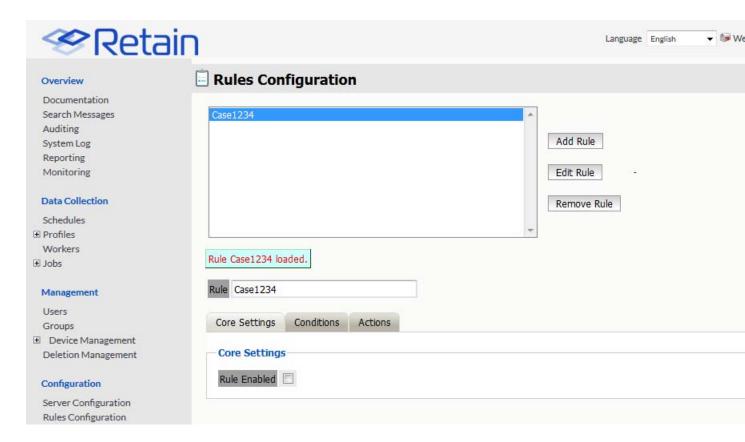
The REST API is an input protocol only. For Archive security, the API prohibits migrating data out of the Retain Archive. There are also limits placed in each applications API license file to protect the integrity of the data being input into the archive, based on the application.

To gain a credentialed key for a third party application, contact a GWAVA sales representative. Use of the REST API will require a REST API base license, and an application specific license for each desired application.

Rules Configuration

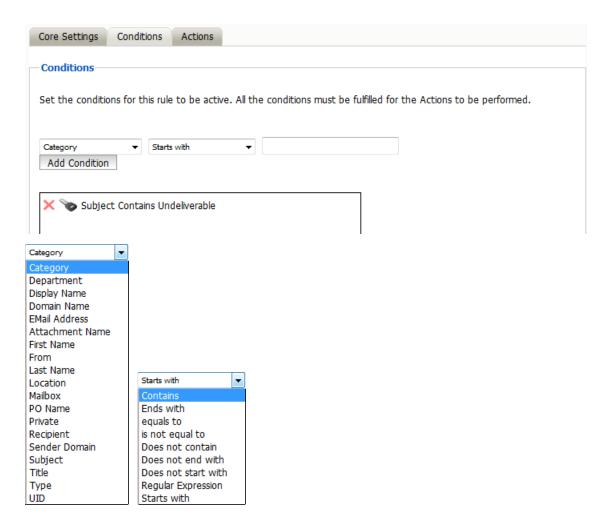
Rules for Retain is a way for administrators to automate tasks and administration based on criteria. This means that any message containing items or information specified in the criteria to have the described action applied automatically. A rule must be enabled, have at least one condition, and one action defined before a rule comes into effect.

To create a rule, select the 'add rule' button and name the rule. To make the rule active, select 'Rule Enabled".



Conditions

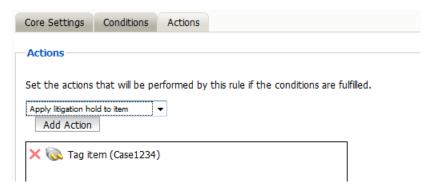
The conditions are the criteria which determine whether the action is applied to message data or not.



Criteria is entered in plain text or in regular expression, and are separated into categories. Select the category that the criteria is to be restricted to, from the menu.

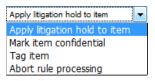
Once a category has been selected, the criteria to be applied to that needs to be specified. The 'Starts with' menu further defines the criteria. The 'Starts with' menu also allows the administrator to either relax some constraints by selecting a vague entry like 'Contains' or 'Starts with', or specify the criteria with 'Regular Expression'. Select the appropriate limiter and then select the 'Add Condition' button. Multiple conditions may be added. Make sure to save changes before leaving the page.

Actions



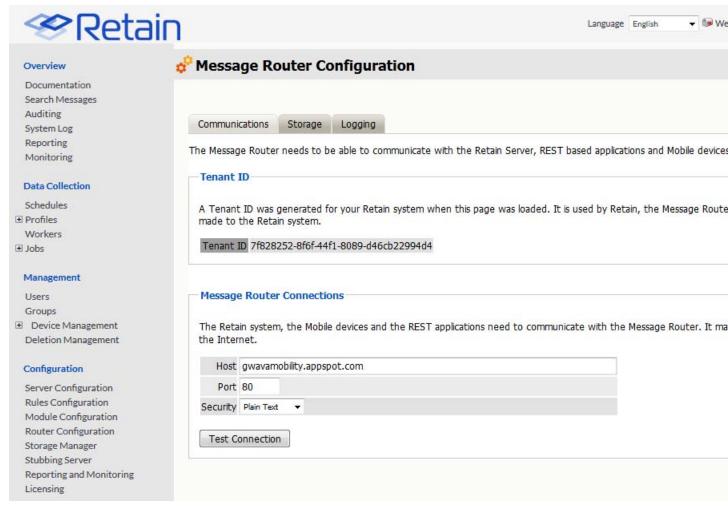
The actions tab defines which actions are to be taken. There are four different options to choose from.

All four actions may be applied to any one rule. Once the desired actions are applied to a rule, save changes before leaving the page.



Router Configuration

Configuration for the Message Router may be changed or checked on this page. The Router is used in conjunction with the mobile module, the BBM module integration, and any other REST API integrations. In contrast with the Retain Server, the Router is best placed in a network's DMZ, enabling it to be fully accessible to communication from the Internet as well as the network.



The Router's communication settings and port for both the Retain Server and the Internet is kept here. This is the address that the mobile devices will use to talk to the router and should be accessible to the outside world.

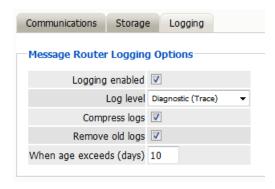
Storage



The storage tab contains the path to the temporary storage location local to the REST Router. The storage is used by the REST Router to store data for a few seconds before it is sent on to the Retain Server. In the event of a failure to access the Retain Server, the REST Router will house the data until the Server is accessed. This storage area is not required to be large, and a few gigabytes is more than sufficient.

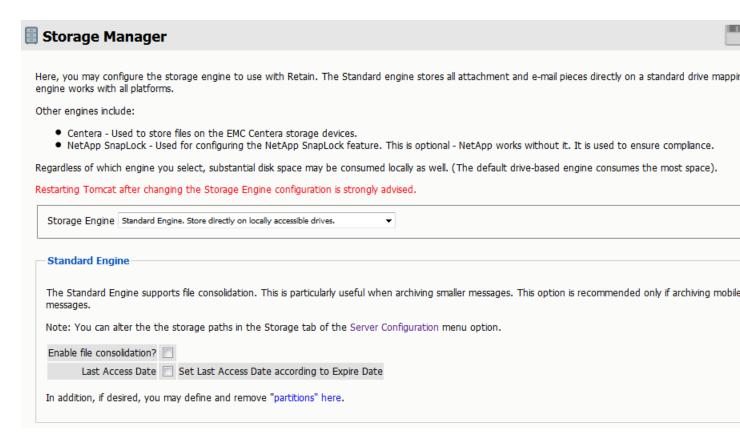
Logging

The logging tab contains settings to enable logging, dictate the logging level of the Router, and settings to purge old logs from the system after a specified age.



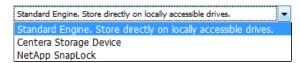
Storage Engine

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 two options are:

- Centera Storage Device
- NetApp SnapLock

Enhanced Standard Engine

The Enhanced standard storage engine is the standard storage option which should be selected for all new installations of Retain. The storage engine utilizes locally accessible drives and media. When selecting this option, make sure that storage requirements will be easily met, size or capacity increased, and managed.

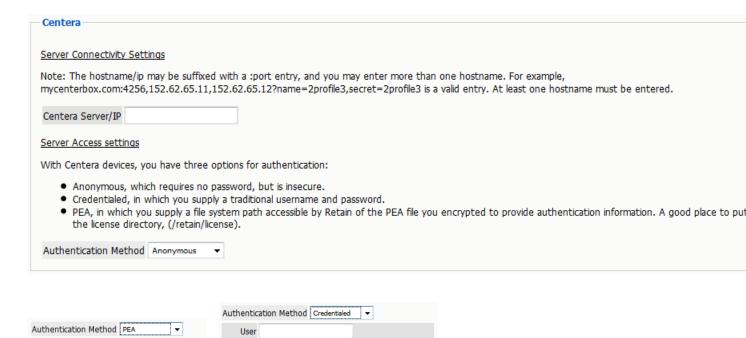
The enhanced engine also allows for background compression to be used. Background compression is recommended for social media and mobile modules, to help keep the size of the data store down.

Legacy Engine

The legacy engine works as a seamless connection to the 2.x line of Retain, and is available as a bridge for upgraded devices and legacy support systems.

Centera Storage Device

The Centera Storage option simply requires you to put the storage server IP address or DNS name, and the authentication method.



(Password is not set)

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.

Password

NetApp SnapLock

PEA Path (full path)

NetApp is fully supported with SnapLock for storage. Note the information in the option. To enforce WORM rules you MUST set the Job Expiration feature. Also set the IP address, user name, password, and the correct path to the server. The path must be set in netapp format.

─ NetApp SnapLock	
— несарр энарсоск	
<u>SnapLock settings</u>	
	e with a Job specifying an expiry date, their retention date is set according to your Jo etting, if preferred, in which case none of the parameters on this screen need to be
	IC path. There is no definite correlation for example between this path and a share path and the Retain blob files. Consult your NetApp Technical Support for more information on
Enable SnapLock? № ▼	
NetApp Server/IP	
Port 80	
Admin User	
Admin Password	(Password is not set)
Base path to archive directory (example: /vol/myvolname/archive)	

Partitions

Retain also supports the ability to define a different storage location, through the partitioning option. Partitioning is only available when Storage is set to the 'local' drive. To access the partitions option click on "partitions" in the Enhanced or Standard Engine window. Ideally, this option would only be used to archive to a new, larger, location when the current storage path has been filled, or to separate mail storage between years. (I.e. /archive/2008; /archive/2009)

- Partitions
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 prefe 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 da 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: /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
New Partition Path

A Retain system without a specified partition location will store all mail in the default storage location. This is defined in the Server Configuration menu, under the Storage tab mentioned previously. The default location is set during the initial install.

Read the bold warning on changing partitions!

You may only specify one new partition at a time. After you have entered the new partition information, click 'save changes'. *Creation or changes to a partition requires Tomcat to be restarted.* When adding a Partition on Linux, the path needs to be created manually on the Linux Server before specifying the path within Retain. Ensure that the permissions for the new path are given to Tomcat.

Once Tomcat has been restarted, the new partition can be viewed, edited, or deleted from this window.

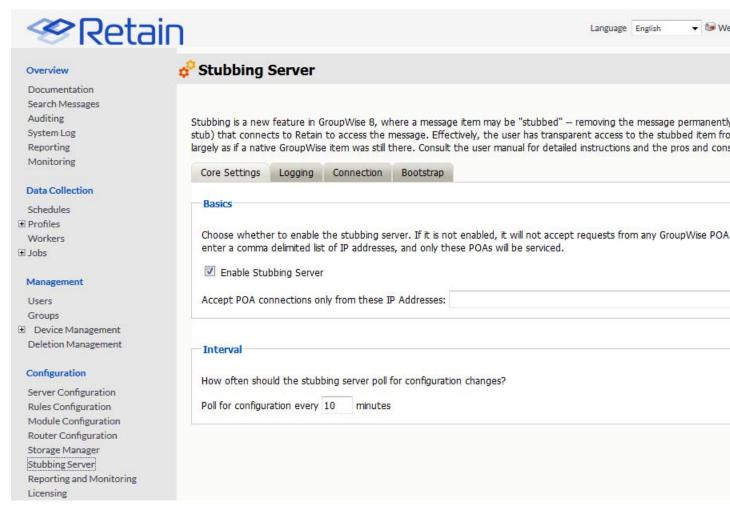
Partitions are listed with their Begin Date, Name, and Path. The Begin Date is the creation date. A partition is used to store all archived mail from the moment it is created

Because of how Retain indexes and tracks messages that are archived, deleting a partition can lead to catastrophic data loss. Do not delete an active storage partition or location without first consulting support.

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

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)

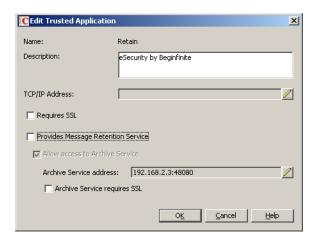
Stubbing Server



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

- 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.
- ConsoleOne must be configured as discussed below.
- 3. Retain job profiles must have Stubbing turned on. (see the Profiles Section)

Stubbing server must be installed and setup as discussed below. (For install see your platform install guide.)

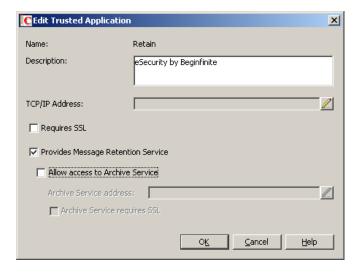


ConsoleOne GroupWise settings

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

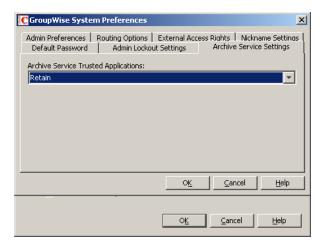
- Connect to the domain and highlight the domain or GroupWise system
- Select Tools | GroupWise System Operations | Trusted Applications
- Select Retain and click Edit
- Select "Allow Access to Archive Server"
- Put the address or DNS name of the Stub Server and the correct port as shown (48080)
- Click OK and exit the rest of the windows

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



There also needs to be a change in the settings in the preferences menu:

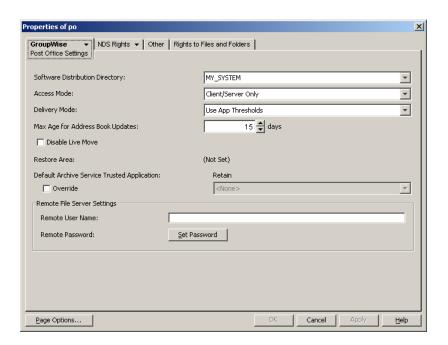
- Connect to the domain and highlight the domain or GroupWise system
- Select Tools | GroupWise System Operations | System Preferences
- Ensure that Retain is set under the Archive Service Settings as a trusted Application. Select it if it is not already selected



Set the Post Office Remote File Server Settings.

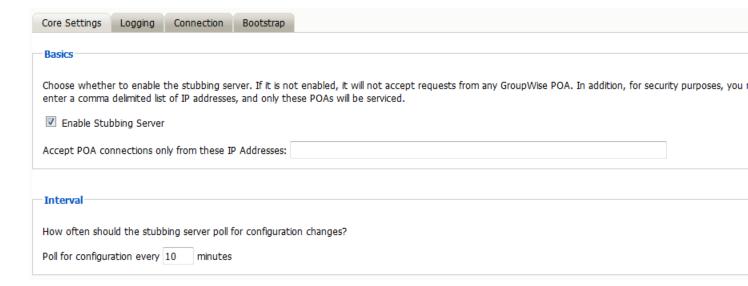
Enter the properties of the Post office. The settings should look as they do here.

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

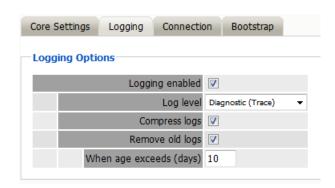


Retain Stubbing Settings

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



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.



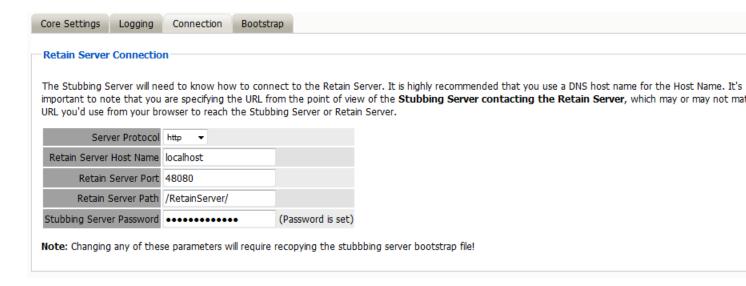
Save changes before continuing.

Logging

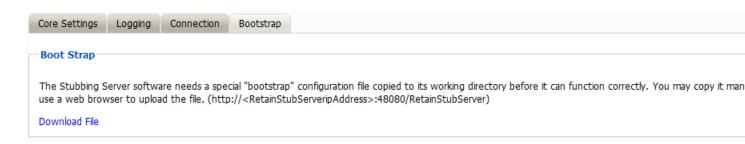
Log levels higher than Normal are not necessary unless troubleshooting or working with Support. It is strongly recommended to compress and remove logs after a time period to protect system space. Set time to desired level.

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).



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



Download and save the bootstrap file.

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

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



Retain Stubbing Server

The Stubbing Server needs to be configured in the Retain Server web interface. Once it has been configured a "bootstrap" file needs to be downloaded and applied here.

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

noda bootsti	ap Configuration		
Browse	lo file selected.		
Upload File			

More Information

To download the bootstrap file and upload it here:

- · Open the Retain server's web interface.
- · Select the Stubbing Server
- 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.

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'.

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.

Continue

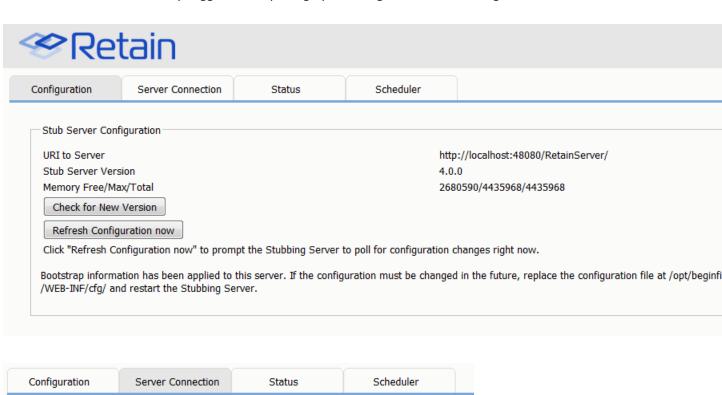


We can now Login to the Stub Server.

Server Connection

Successfully connected to Retain server

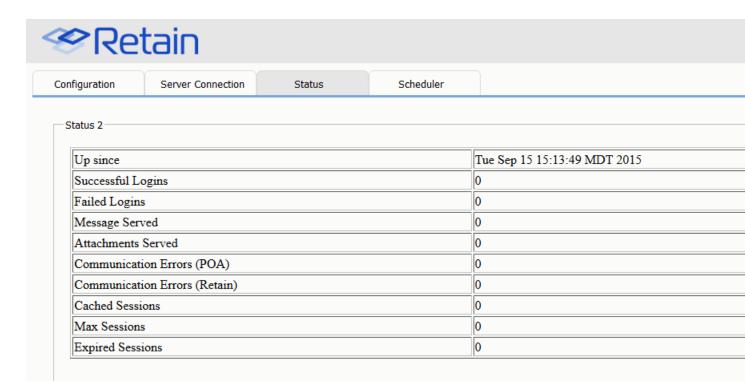
You must Login with a user that has sufficient rights. See the User Rights section to add or remove rights. The Stub Server Interface is nearly identical to the Worker interface. You may check for updated versions of the Stub Server or manually trigger server polling by selecting the Refresh 'Configuration now' button.



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

Status

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



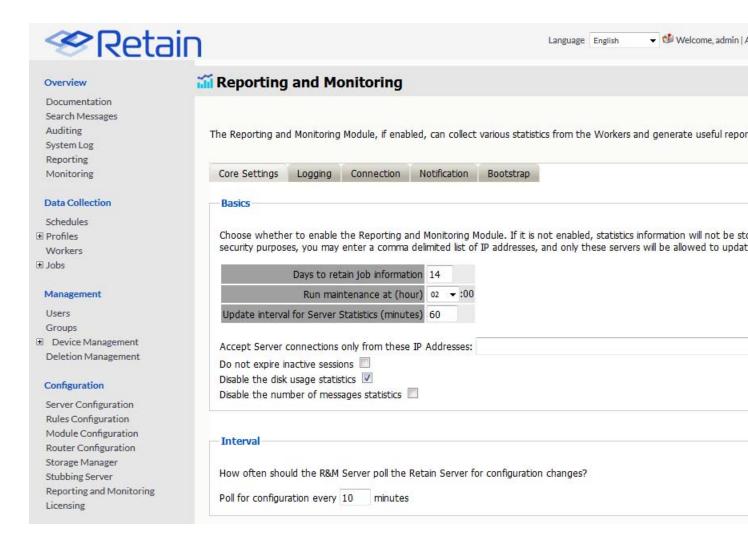


Scheduler

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

Reporting and Monitoring Configuration

Retain comes equipped with a separate Reporting and monitoring server. Reporting and Monitoring may be installed on the same or a different physical or virtual machine as the Retain Server.



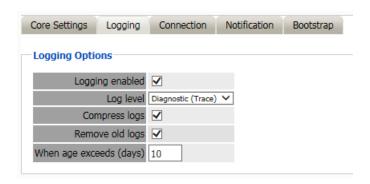
Reporting and Monitoring functions much like the Retain Worker in that it is installed as a separate entity and requires connection information from the Retain Server. Additional security is provided through the option to specify exactly which IP address(es) the Retain Server will allow a connection from. Separate with commas the IP addresses allowed if more than one. No mail data is transferred through this connection, only server statistic information.

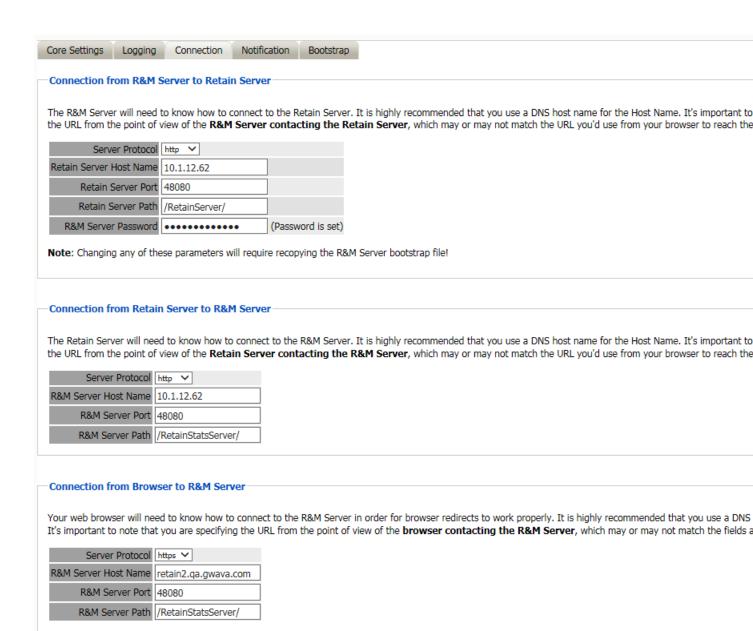
The Core Settings, Logging, Connection and Notification sections of Reporting and Monitoring configuration work identically to the Worker connection settings, and should be filled-out accordingly. Most server implementations will require little to no modification to these settings save the notification tab. Limiting the server connections to specified IP Addresses limits which machines have access to the Reporting and Monitoring module. Separate multiple addresses with a comma.

Putting a checkmark in the 'Do not expire inactive sessions' will allow sessions with the R&M module to never expire. This is useful in conjunction with the 'monitoring' page.

Selecting the 'Disable the disk usage statistics' box will disables the 'archive size' from server utilization report. Selecting 'Disable the number of messages statistics' disables the 'total messages stored' on the Reporting and Monitoring dashboard. These options are for systems which encounter a bottle-neck which impacts system performance with these checks.

Set the logging level to the desired setting. Leaving the logging as default 'normal' is recommended unless instructed to increase the logging level by Support.

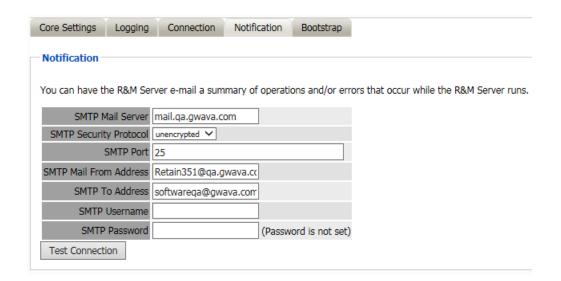




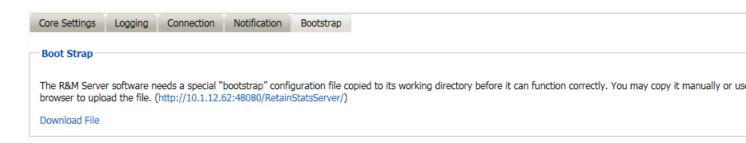
The Connection between Reporting and Monitoring, and the Retain Server should not be changed unless the Retain Server requires a modified connection path, (such as a custom port.)

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

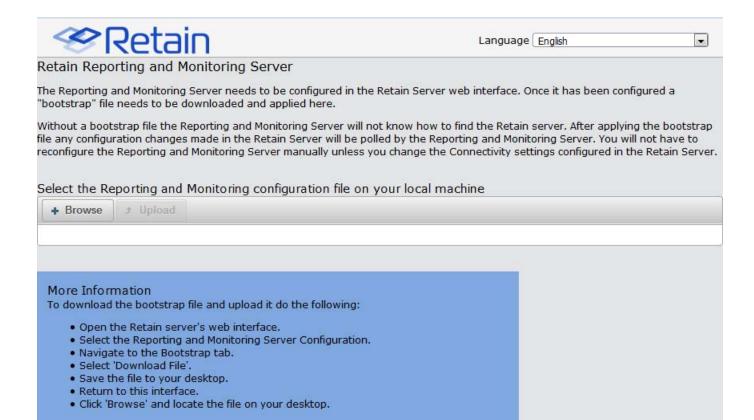
The connection from the Browser to the R&M Server is how the browser is going to connect to the R&M server when redirected from links in the Server. This connection information needs to be accurate for where the R&M Server is installed and listening.



Notification defaults to the standard notification address specified in the Retain Server activation and initial configuration. If the notification is desired to be different, modify the settings here. Make sure to test the connection to ensure proper function.



Like the Retain Worker, the R&M server must be uploaded a saved configuration or 'bootstrap' file. To access this tab, the configuration must be saved first. Then the bootstrap tab will become available. Download the file to a known location and then browse to the R&M server webpage. The default connection address is shown on the connection tab.



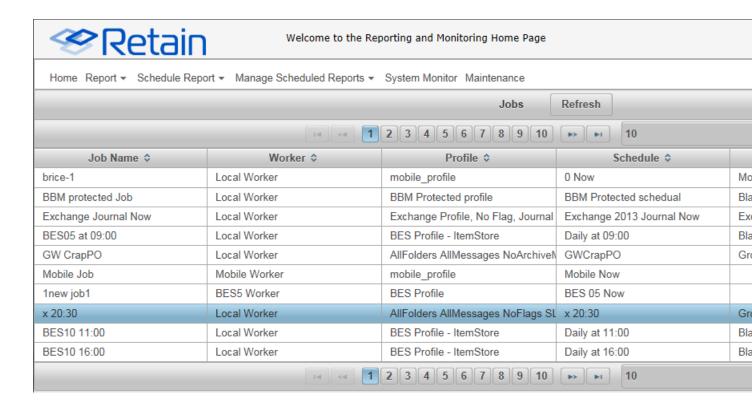
From the R&M initial webpage, click on the "Browse" button and browse to the saved bootstrap file and select it.

Click 'Continue' to upload the bootstrap file. As the bootstrap file is uploaded, R&M functions will become available. Only users with the administrative "Access Reporting and Monitoring Server" right will be able to login to the stat server.

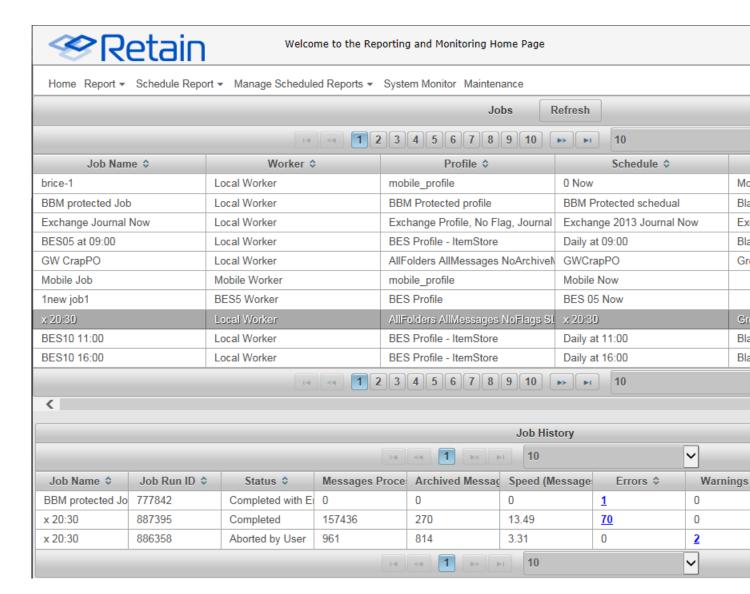
To log into the R&M server, select the 'Reporting' or 'Monitoring' links from the Retain Server left menu to spawn a new tab or page to the R&M module.

Home

On login, Reporting and Monitoring opens to the Home page. The Home page is also accessible through the 'Home' button along the top bar. The R&M server has different page links and menus across the top of the menu bar, accessing different functionality. The R&M server has many tables and columns displaying a large amount of information. To make this information more accessible, every column has been made sortable. To sort the columns, simply click on them. They may be sorted to show the highest number first, lowest number first, or not sorted at all.



The Home page list the different jobs, associated workers, profiles, schedules, and last reported run times configured for the system. These jobs can be selected to display the job history and detailed information for each instance of when that job was run.



Further selection of a listed job instance will display details for that selected run of the selected job. If an error exists on any job, it will show up as a linked number under the error column. Errors reported here are linked to information on the error which has occurred.

A selected error will spawn a new window which displays the error and any applicable information, such as any associated files or error code numbers, a quick description (if a Retain error), the severity level, how often it has occurred, and in how many different mailboxes it has occurred.

The last error encountered in a job is also displayed in the Job details window. Select a column header to sort it as desired. The Errors column shown below has been sorted to show the highest number of errors first.

			Job Details
Job Name	x 20:30	Status	Completed
Worker	Local Worker	Duration (hh:mm:ss)	03:18:59
Start Time	2015-09-17 20:39	End Time	2015-09-17 23:57
Mailboxes Processed	238	Archived Messages	270
Messages Processed	157436	Message Data Stored	2.14 MB
Last PO Scanned	gw_po	Last Mailbox Scanned	mkures
Errors	70	Last Error	java.io.IOException: GWPOA HTTP server

	Job Mailbox Summary						
		14 44 1 7	2 3 4 5 6 7	8 9 10 🕟 ы	10		
Mailbox \$	Messages Processed	Archived Messages	Errors >	Warnings ≎	Advanced Date Flags	Start	
mkures(gw_po)	26	0	0	0	Last Archived Flag (2015-08-26 07:18)	2015-09	
mgondek(gw_po)	27	0	0	0	Last Archived Flag (2015-08-26 07:18)	2015-09	
BradH(IS&T)	47	0	0	0	Last Archived Flag (2015-08-26 15:18)	2015-09	

Report and Schedule Report

Report and Schedule Report functions are accessed through drop-down menus in the interface. To access the different reports, select the 'Report' or 'Schedule Report' menus from the top menu and then click on the desired report. Most reports which are accessible through the 'Report' menu are also available as a 'Scheduled Report', except for 'Job Run Mailbox Summary'. All reports under the 'Report' menu except 'Job Run Mailbox Summary', can be scheduled and have a 'Schedule' button at the bottom of the delivered HTML report. Selecting this button copies the settings used to generate the visible report to the Schedule Report interface. This convenience is designed to allow rapid fine-tuning of reports to deliver the desired information. First select the report you wish to run, and then fine-tune the report until it reaches the desired effect. Once the desired settings have been achieved, select the 'Schedule' button from the bottom of the report to carry all settings into the corresponding 'Schedule Report' interface, and schedule the report.

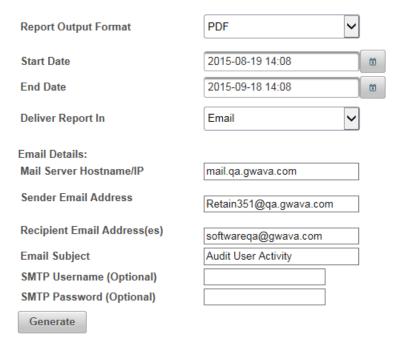
Report *	Schedule Report *			
Audit User Activity	Audit User Activity			
Export Audit Table	Export Audit Table			
Job Run Averages	Job Run Averages			
Job Run Summary	Job Summary			
Job Summary	Archive Summary by			
Archive Summary by User	Server Utilization			
Server Utilization	Job Mailbox Summary			
Job Run Mailbox Summary				
Job Mailbox Summary				

Reports differ from Scheduled reports in more than immediate delivery, reports support HTML in addition to PDF and XLS delivery formats. However, both report structures may be emailed to a desired recipient. When a report is created and emailed it is formatted as selected and then attached to a message which is sent to the desired recipient.

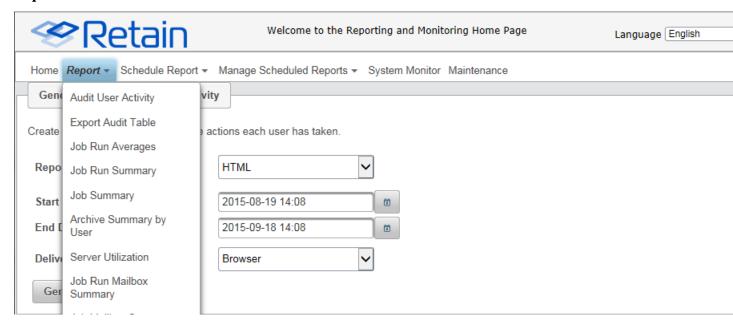
When the Email delivery option is selected, additional configuration is required. The R&M module needs to know not only what address to send the report to and what mail server to use, but also allows the configuration of the email subject and sender address. If an SMTP username and password are required for the desired host system, they should be provided.

The R&M module will automatically copy the SMTP configuration from the Retain Server. If this information is not correct, it should be configured here before selecting the 'Generate' button.

This configuration is the same interface for all emailed reports, immediate and scheduled.



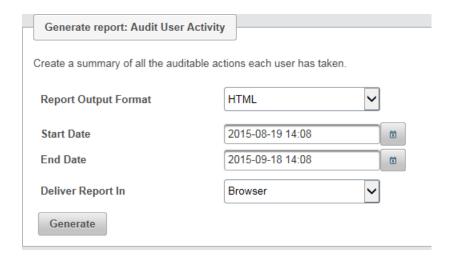
Reports

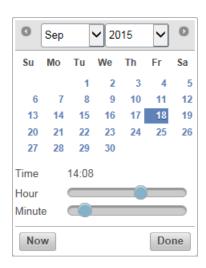


Reports may be created in HTML, PDF, or XLS, and may be delivered straight to the Browser or in a specified email. Different reports have different settings, however the basics are the same; select or specify the desired criteria, specify the format and delivery method, then click on the "Generate" button to create the report.

Audit User Activity

The Audit User Activity report details the different actions performed in the Retain system by different users. This report is a detailed list of all auditable actions taken and which users took them in the Retain system.





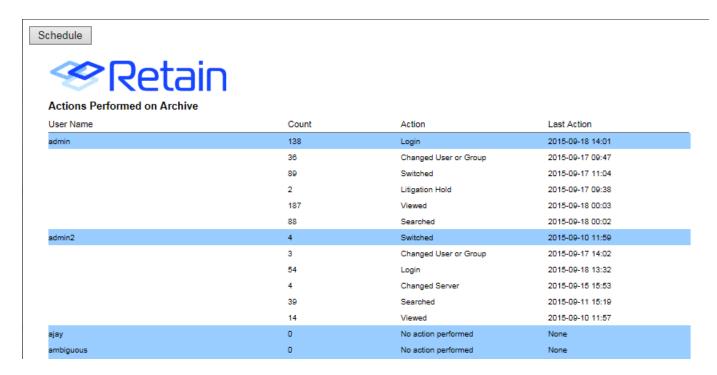
While this report lists all actions by all users, it may still be restricted by a date range. The date ranges are specified through a calendar interface.

The date and year are selected through the top bar, while the hour and minute, if desired, are selected through sliders. The date can be immediately set to the current moment by selecting the 'Now' button. When configuration has been completed, select the 'Done' button to set the date. Both the Start and End Date have identical date tools.

When the settings are as desired select the 'Generate' button to create the report.

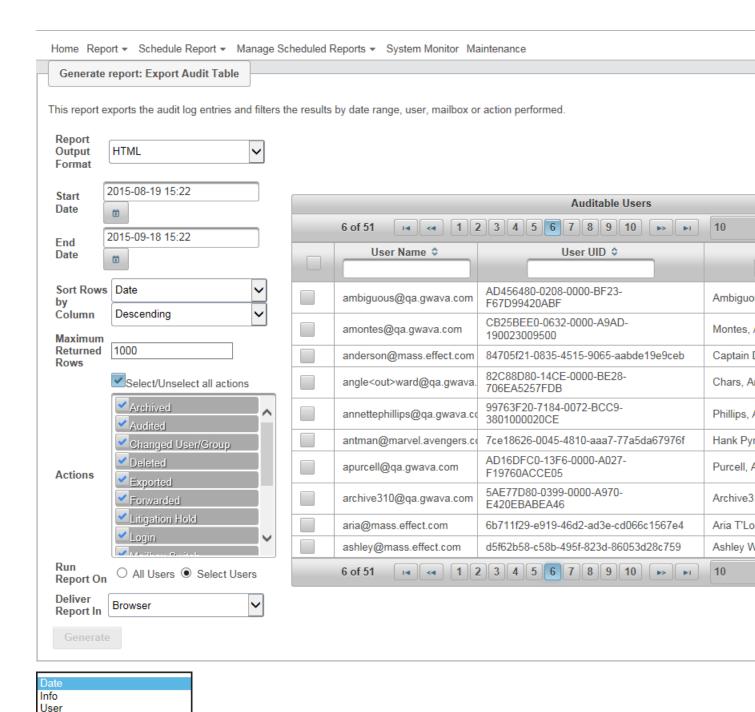
The report lists all users, a quick description of the action taken and how many times it was performed, and the date of the last time it was performed. Actions such as viewing an item, searching the archive, changing settings on either the server or in the interface, forwarding items, applying litigation holds, pub-

lishing items, etc. To take the same settings used when creating this report, select the 'Schedule' button to enter the Schedule Report interface with the same settings.



Export Audit Table

The Export Audit Table sorts the entire Audit log with the settings desired and creates a report for specific users or all users, with specific events or all events. This report can be limited in size and date.



If restricting the report to a specific user or users, a menu displaying 'auditable users' is displayed. The different fields are searchable. The option to select all, or none, is also present. Users may be searched for via User Name or email address, User UID, or the user Display Name.

Mailbox Job Name Action

The rows and columns are not sortable in the generated report, but they can be sorted or organized from the report generation interface. The rows may be organized by the columns listed. This has the exact

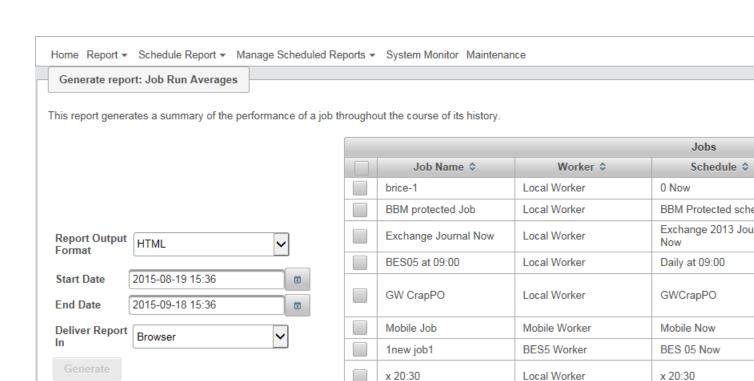
same effect as organizing columns for the jobs on the reporting and monitoring home page. Configure as desired and then select 'Generate' to create the report.



To Schedule the report, select the 'Schedule' button from the top of the report to keep the settings.

Job Run Averages

The job run averages report takes the performance of any selected particular job and creates a chart to display the trends. The effectiveness, overall load, and performance can be rapidly displayed. The Job Run Averages report requires a single job to be selected. The 'Jobs' selection window has columns which are sortable for convenience.



BES10 11:00

BES10 16:00

1 of 11

Local Worker

Local Worker

Daily at 11:00

Daily at 16:00

1 2 3 4 5 6 7 8 9 10



Once the date period has been configured, select 'Generate' to create the report.

Performance trends are immediately apparent as run time is compared with message counts over a time period. The averages used to create the graph are listed above the graph, as well as a legend below.

As with the rest of the reports which exist in the 'Schedule Report' interface, to schedule this report with the settings which created it, select the 'Schedule' button from the top of the report to use the same settings and be linked directly to the 'Schedule Reports' interface.

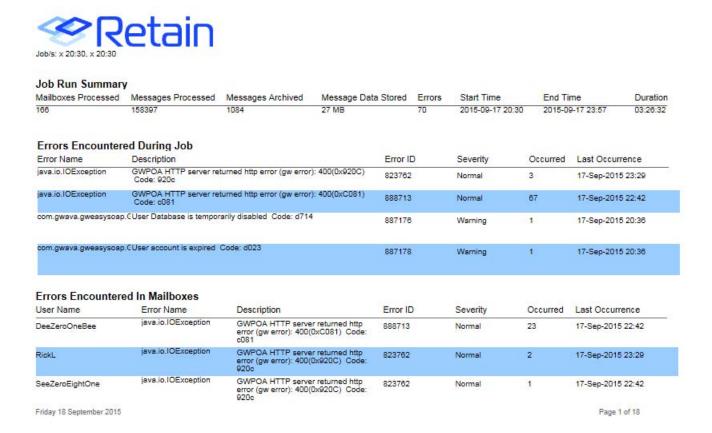
Job Run Summary

The Job Run Summary report gives a complete report of the performance of a single job. Whereas the Job Run Averages report charts raw performance averages, the Job Run Summary report goes into greater detail on the total number of archived messages, the errors, and mailbox summary for a single instance of a job.

Job Run Summary and Job Summary reports have the option to limit the number of reported errors displayed.

Generate report: Job Run Sun	nmary			
Concrete reports out real out	ary			
This report generates a summary f	or a single job run that includes the number of messages archived	, the en	rors that have occurred, and	a summary by
Report Output Format	HTML		Job Name	
Maximum Error Item Occurrences	1000		x 20:30	Start
Occurrences		•	x 20:30	2015-09-17
Deliver Report In	Browser		x 20:30	2015-09-17
Generate			1 of 1	<4 1

Select the desired single instance of a job and then select 'Generate'.

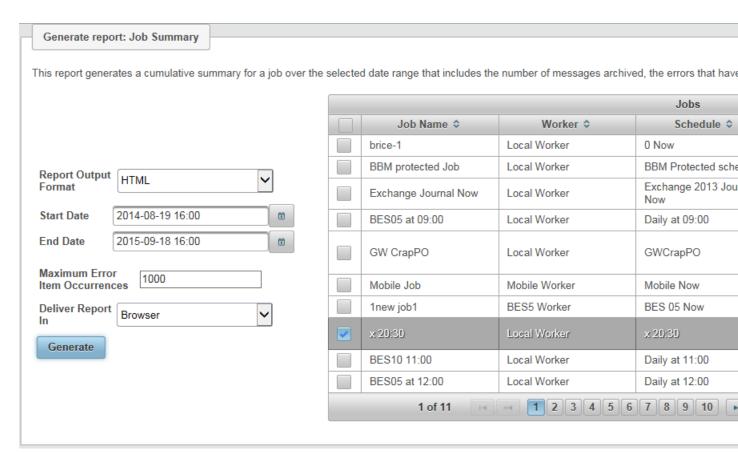


The summary displays the total mailboxes, messages, data store, errors, and running time of the job. Afterwards, this information is broken down by each mailbox the job covers.

Job Summary

The Job Summary report is much the same to the Job Run Summary report, except that this report covers the entire history of a job and not just a single run. To run the report, select the desired job and any time period, if desired, the desired output and delivery method, and click 'Generate'.

Job Summary contains the option to limit the number of reported errors displayed.



The report generated displays the number of mailboxes that the job processes, the total messages the job has processed, how many messages are actually archived, the message data stored, the total errors encountered, when the job first started, and the total time the job has been active.

Immediately below the initial information, a pie chart displays the amount of messages actually archived in comparison to the number of messages the job processed.

Finally, below the pie chart is a summary, by individual mailbox, of job actions. Each mailbox is displayed with the pertinent information, including the messages archived, messages processed, critical errors, a record of the advancement of message store flags, and the duration of the job in that mailbox.

To schedule this job with the settings desired, select the 'Schedule' button from the top of the report. All settings used to create the report will automatically be copied into the 'Schedule Report' interface.



			-				
- 1	-	h	•	m	m	•	rγ

Mailboxes Processed	Messages Processed	Messages Archived	Message Data Stored	Errors	Start Time	End Time	Duration
166	158397	1084	27 MB	70	2015-09-17 20:30	2015-09-17 23:57	03:26:32

Errors Encountered During Job

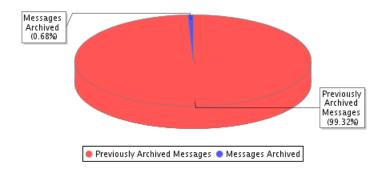
Error Name	Description	Error ID	Severity	Occurred	Last Occurrence
java.io.IOException	GWPOA HTTP server returned http error (gw error): 400 (0x920C) Code: 920c	823762	Normal	3	17-Sep-2015 23:29
java.io.IOException	GWPOA HTTP server returned http error (gw error): 400 (0xC081) Code: c081	888713	Normal	87	17-Sep-2015 22:42

Errors Encountered In Mailboxes

User Name	Error Name	Description	Error ID	Severity	Occurred	Last Occurrence
DeeZeroOneBee	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0xC081) Code: c081	888713	Normal	23	17-Sep-2015 22:42
RickL	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0x920C) Code: 920c	823762	Normal	2	17-Sep-2015 23:29

SeeZeroEightOne	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0x920C) Code: 920c	823762	Normal	1	17-Sep-2015 22:42
SeeZeroEightOne	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0xC081)	888713	Normal	44	17-Sep-2015 22:42
Eriday 18 Sentember 2015						Page 1 of 19

Archive Ratio of Processed Messages



Friday 18 September 2015 Page 6 of 18

Job/s: x 20:30

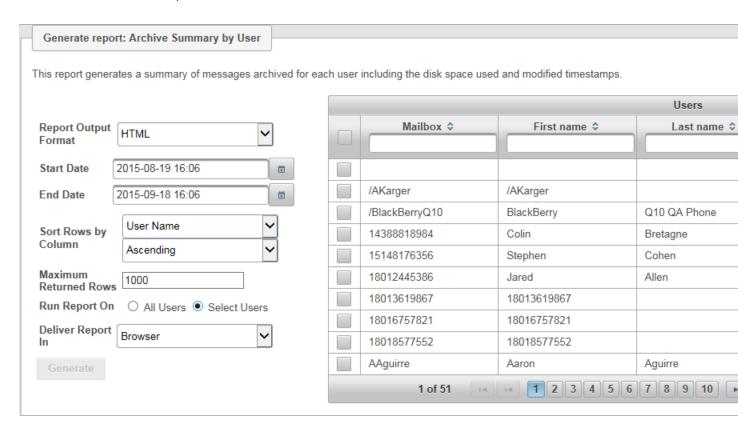
Archive Summary per Mailbox

Mailbox	Processed	Archived	Errors	Warnings	Advanced Date Flags	Start Time	End Time
88	218	0	D	D	Last Archived Flag (2015 09-08 17:07)	- 2015-09-17 22:40	2015-09-17 23:42

Archive Summary by User

The Archive Summary by User report creates a summary of all users', or a specific user's, archive history. The summary includes all pertinent information for each mailbox's archive. To run the report, either all or one mailbox must be selected. If all mailboxes are selected, the maximum returned rows may be a very important setting to keep the report from becoming immense.

While the report cannot be sorted after it is generated, it may be sorted at generation through the 'Sort Rows by Column' feature. The columns may be sorted by either descending or ascending order. Select 'Generate' to create the report.



The report displays the total information for every mailbox selected; User name, the time stamps and their advanced dates, and the message store for the associated user.





Total Users in Address Book: 10367

Messac	ies Ar	chived	l by	User

a0conn01			Timestamps	
	Archived:	9	Item Store Flag: Not set	
(EXCHANGEQA)	Processed:	83	Last Archive Flag: 2015-07-30 08:05	
	Disk Space Used:	0 MB	Retention Flag: Not set	
Data Since: 2015-06-05 02:00			Purge Flag: Not set	
a0izyu01	Archived:	0	Item Store Flag: Not set	
(EXCHANGEQA)	Processed:	0	Last Archive Flag: Not set	
	Disk Space Used:	0 MB	Retention Flag: Not set	
Data Since: 2015-07-24 09:04			Purge Flag: Not set	
a0john01	Archived:	8	Item Store Flag: Not set	
(EXCHANGEQA)	Processed:	35	Last Archive Flag: 2015-07-30 08:05	
	Disk Space Used:	0 MB	Retention Flag: Not set	
Data Since: 2015-06-21 02:00			Purge Flag: Not set	
aa	Archived:	2	Item Store Flag: Not set	
(LXPO)	Processed:	679	Last Archive Flag: 2015-09-08 17:07	
	Disk Space Used:	0 MB	Retention Flag: Not set	
Data Since: 2014-02-10 21:33			Purge Flag: Not set	
Administrator	Archived:	1	Item Store Flag: Not set	
(LXPO)	Processed:	454	Last Archive Flag: 2015-08-26 15:18	
	Disk Space Used:	0 MB	Retention Flag: Not set	
Data Since: 2014-06-14 02:00			Purge Flag: Not set	

Friday 18 September 2015 Page 1 of 35

The most valuable part of this report, and what will make this report a staple for administrators, is the ability to monitor and get a concise report for the disk space used for each individual user along with the messages processed and stored.

If the report is to be scheduled, select the 'schedule' button from the top of the report to use the same settings which created the report.

Server Utilization

The Server Utilization report is a direct report on the load that the Retain Server is under. The only option is to limit the date range for which the report is active. Select the report period and then select 'Generate' to produce the report.

Generate report: Server Utilization		
This report summarizes the server healt	h as it changes over time. This	s report includes information such as CPU and memory usage, index and archiv
Report Output Format	HTML	<u>~</u>
Start Date	2015-08-19 16:11	
End Date	2015-09-18 16:11	
Deliver Report In	Browser	<u>~</u>
Generate		

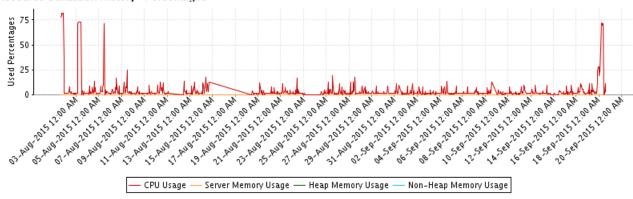
The Server utilization report details the memory, CPU utilization, Archive size, Thread usage, Index size, heap memory, and non-heap memory usage all in quick to digest charts.

To schedule the report, select the 'Schedule' button from the top of the report to utilize the same settings.



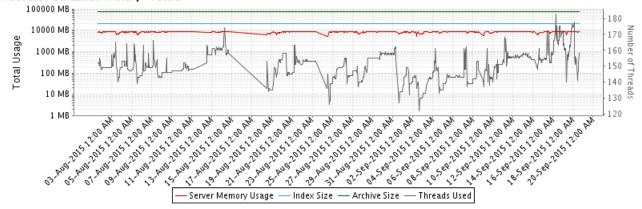
	Average Used	Minimum Used	Maximum Used	Total Available
CPU Usage	4.57 %	0.00 %	82.00 %	Not Applicable
Server Memory	8,273.98 MB	4,900.70 MB	8,611.11 MB	8,731.74 MB
JVM Heap Memory	2,139.97 MB	642.09 MB	4,019.47 MB	4,326.50 MB
JVM Non-Heap Memory	1,882.90 MB	164.05 MB	4,735.27 MB	0.00 MB
Index Size	18.41 GB	16.43 GB	21.38 GB	Not Applicable
Threads	149.93 GB	122.00 GB	183.00 GB	Not Applicable
Archive Size	72.46 GB	72.36 GB	72.57 GB	Not Applicable

Resource Utilization History - Percentages



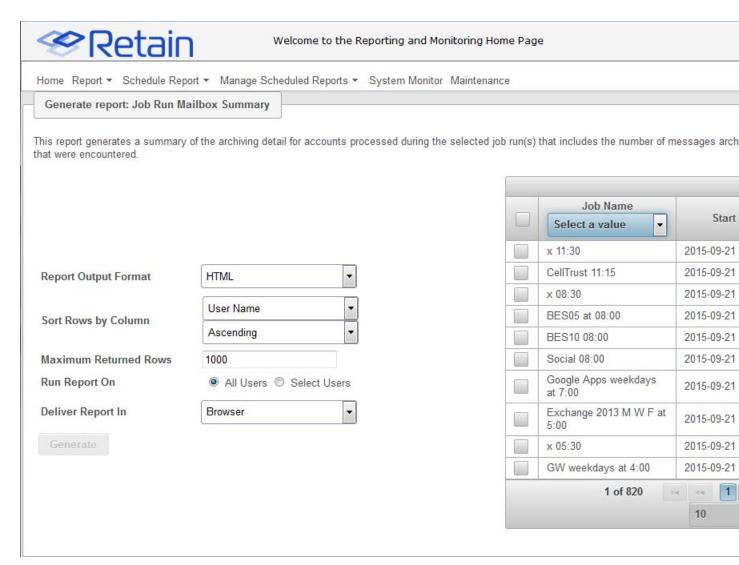
Friday 18 September 2015 Page 1 of 2





Job Run Mailbox Summary

The Job Run Mailbox Summary is much like the Job run Summary except it is dedicated to specific mailboxes and specific jobs. This job requires both a job, and a user or users to be selected.



Select the desired job and either all users, or only select users. If desired, the columns may be sorted by a selected value through the 'Sort Rows by Column' drop-down menu. Once set as desired, select the 'Generate' button.

The report shows the details for the selected job and mailboxes. The user name, messages archived, disk space used, and archiving time is all displayed for each mailbox as well as for the total job.



Archive Report Totals

 Messages Archived
 122

 Disk Space Used
 0 MB

 Archiving Time
 00:12:54

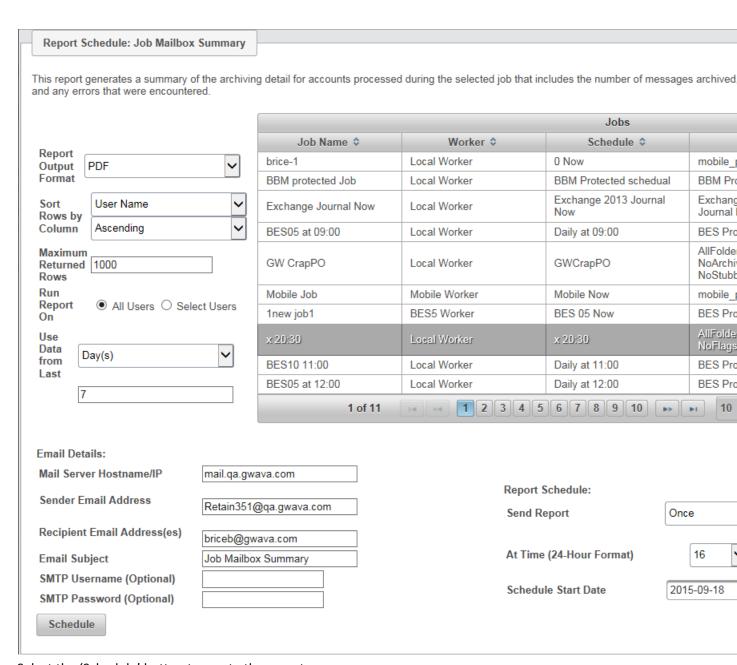
Archiving Detail by User

User Name	Messages Archived	Disk Space Used	Archiving Time	
Rasoli, Hakim	111	0 MB	00:00:18	
User, Stubby	11	0 MB	00:06:21	
User, Stubby 14	0	0 MB	00:06:15	

This report cannot be scheduled.

Job Mailbox Summary

The Job Mailbox Summary is a scheduled report only, that displays the number of messages archived, the disk space used, and the time the job took, separated by separate mailboxes. This report may be restricted to a date range, a limited amount of rows, and may be performed on all, or only select users. This report may also sort the data according to a selected column, in either ascending or descending value.



Select the 'Schedule' button to create the report.



Archive Report Totals

Messages Archived 1084

Disk Space Used 27 MB

Archiving Time 03:19:18

Archiving Detail by User

User Name	Messages Archived	Disk Space Used	Archiving Time
aa	0	0 MB	00:00:18
Administrator	0	0 MB	00:00:04
Allen, Jared	4	0 MB	00:06:59
Alungh, Suman	160	4 MB	00:00:41

Schedule Report

Scheduling reports is accomplished through the Schedule Report interface. Every report which can be scheduled is also available as an instant report with a 'schedule report' button located at the top of the HTML report. Because these instant reports are easier to fine-tune, it is recommended to first configure a scheduled report through the 'Report' interface first, to allow immediate fine tuning and viewing of the results. Once achieved, selecting the 'Schedule' button on the report will take all settings used to generate that report and input them into the schedule Report interface. To schedule any report, a destination email address, SMTP information, and a set schedule are required.

Scheduled reports differ from normal reports in that, due to the scheduled nature, they must be given a time to generate the report and be given an email address to deliver the report to. Reports are delivered as an attachment in an email sent to a specified user. As such, scheduled reports are only offered in PDF and XLS formats.

Greate a summary of all the addita	ble actions each user has taken.		
Report Output Format	PDF		
Use Data from Last	Day(s)		
	7		
Email Details:			
Mail Server Hostname/IP	mail.qa.gwava.com	Report Schedule:	
Sender Email Address	Retain351@qa.gwava.com	Send Report	Daily
Recipient Email Address(es)	softwareqa@gwava.com	At Time (24-Hour Format)	00
Email Subject	Audit User Activity		
SMTP Username (Optional)		Schedule Start Date	2015-09-1
SMTP Password (Optional)		Schedule End Date	



Reports which have a time frame attached to them as part of the criteria will have a floating time window which they may take information from. The configuration allows reports to the use data created in the past number of hours, days, weeks, or months. The number may be set by the user.

Also, the frequency of the report may be configured to be either a single instance or a reoccurring report. The frequency allows for hourly, multiple times on specified hours every day, daily, weekly or on any specified number of days in the week, or at any specified day or days in the month.

In addition, a scheduled start and end date allow the administrator to control exactly when the start and stop dates for any specific report can be. The end date is not required, but is optional.

Manage Scheduled Reports



Once reports have been scheduled, it may be necessary to either cancel or modify the frequency or settings of the report. The Manage Scheduled Reports interface allows jobs to be immediately run, modified, or removed.

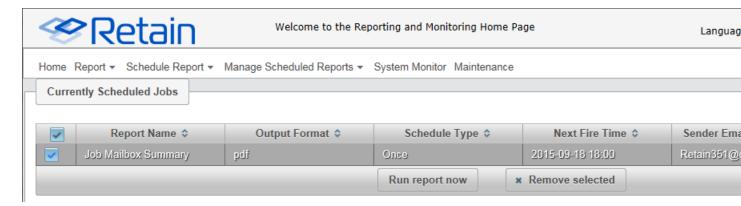
As with all other columns in the R&M interface, the columns are sortable and even searchable to allow quick access to any desired Report through any of the scheduled criteria. For instance, a user which doesn't want to receive a report anymore, but doesn't know what report is being delivered to their email, only needs to know the email subject, or the next delivery time, to immediately identify the report. Unique sender addresses or subjects will readily allow any report to be identified.

To modify or view a report's settings select the 'View/Change details' window.



Once the desired report has been located, select the report and then select the 'View/Change details' button at the bottom to be taken to the Report's configuration. Make any desired changes and select the 'Update' button at the bottom to save the changes.

To remove or immediately run a scheduled report, select the 'Run and Delete reports' window.



Any report or reports may be immediately run or removed from the system through this window. All scheduled reports will be displayed. A select all/none box is at the top left of the table, and allows for a quick selection or clearing of all previous selections which have been made.

To run a report or reports, select the desired report by clicking on the associated checkbox and then select the 'Run report now' button at the bottom. The report will be immediately triggered, and emailed to the specified recipient.

To remove a report or reports, select the desired report(s) by placing a check in the associated checkbox. Once the reports are selected, click the 'Remove selected' button at the bottom of the table.

System Monitor

The System Monitor page is a quick, graphed page which displays the essential stats of the Retain Server system. The R&M module statistics are displayed, with how many reports have been created and of what type. The basic archive progress is also displayed showing the daily, weekly, monthly, and total amount of messages which have been archived in the system.

The system CPU load, memory usage map, and Disk Usage are displayed. The dashboard also contains statistics for disk space saved through single storage instances. These numbers are for the actual host server, and not just the Retain system, with exception of the memory map, which has both. The Memory map is broken out to total system memory free and used, but also includes Java memory statistics — which will be Retain specific unless additional Java applications are running on the Retain Server. This page will not time-out or expire, to allow quick monitoring of the system.

NOTE: A problem has recently been identified with gathering some system-level variables for the System Monitor page. Affected variables include CPU, RAM, and system memory usage. To avoid additional problems these variables have been set to '0' until the issue is resolved.



Welcome to the Reporting and Monitoring Home Page

Language English

Home Report ▼ Schedule Report ▼ Manage Scheduled Reports ▼ System Monitor Maintenance

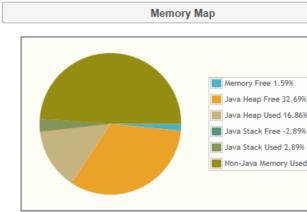
Report Generation Count

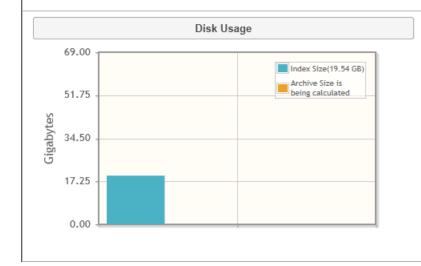
Report Type	Generation Count
Audit User Activity	34
Export Audit Table	12
Job Run Averages	14
Job Run Summary	9
Job Summary	22
Archive Summary by User	11
Server Utilization	10
Job Run Mailbox Summary	27
Job Mailbox Summary	7

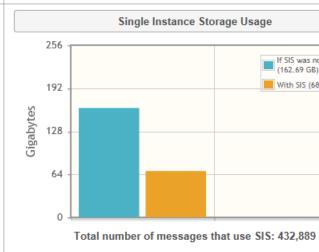
Archiving I	Progress
Messages Archived Today	21
Messages Archived Last 7 Days	4,064
Messages Archived Last 30 Days	16,145
Total Messages Stored	500,000





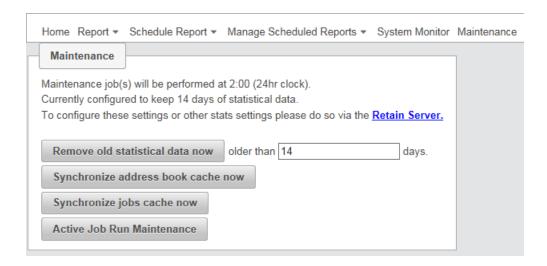






Maintenance

The Maintenance page contains the basic maintenance for the R&M module. This is not maintenance for the entire Retain Server. Because the R&M module maintains individual statistics, it needs to be refreshed and purged at regular intervals. The defaults are shown.



The R&M Module maintains a separate record of the address book and the job cache. The system automatically synchronizes this information with the Retain Server every 24 hours. Still, the different records may become out of sync. If a synchronize job is required before regular maintenance, a manual sync may be individually triggered.

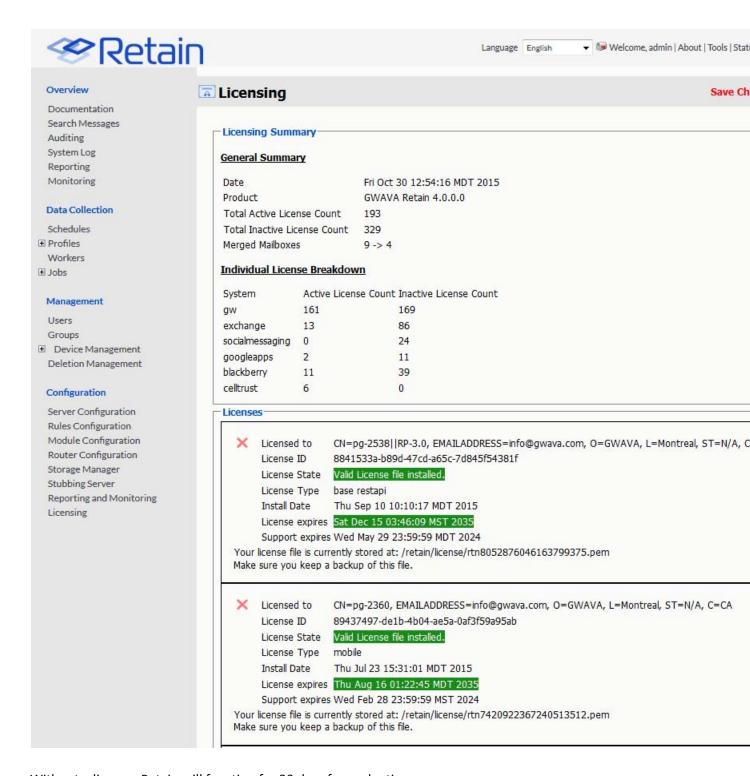
By default, the statistics server only maintains a record of statistics for 14 days, and all data older than 14 days will be automatically expired and the data deleted. If the statistic data is required for longer or shorter, configure and save the setting.

It is not recommended to maintain statistical data older than 30 days.

Licensing

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

Individual licenses which have been already added to the system are displayed under the license summary

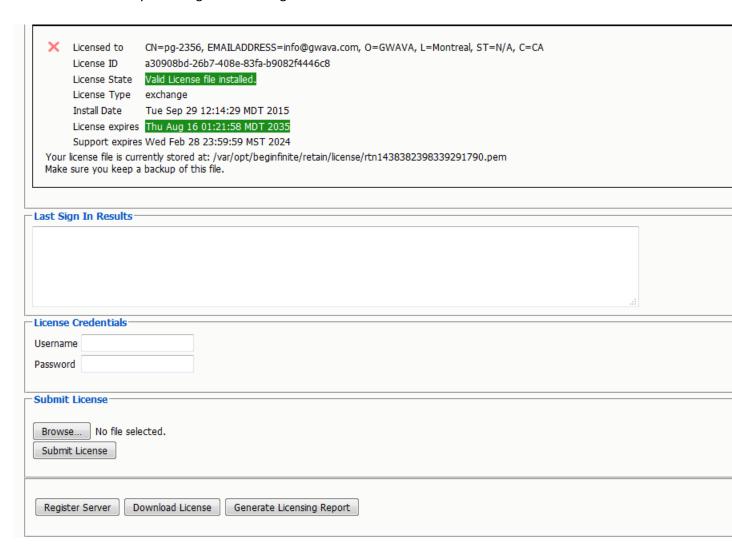


Without a license, Retain will function for 30 days for evaluation purposes.

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



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

A Licensing Report PDF may be generated, which displays the same information as seen in the 'License Summary' section and lists the licenses.

Audit and System Logs

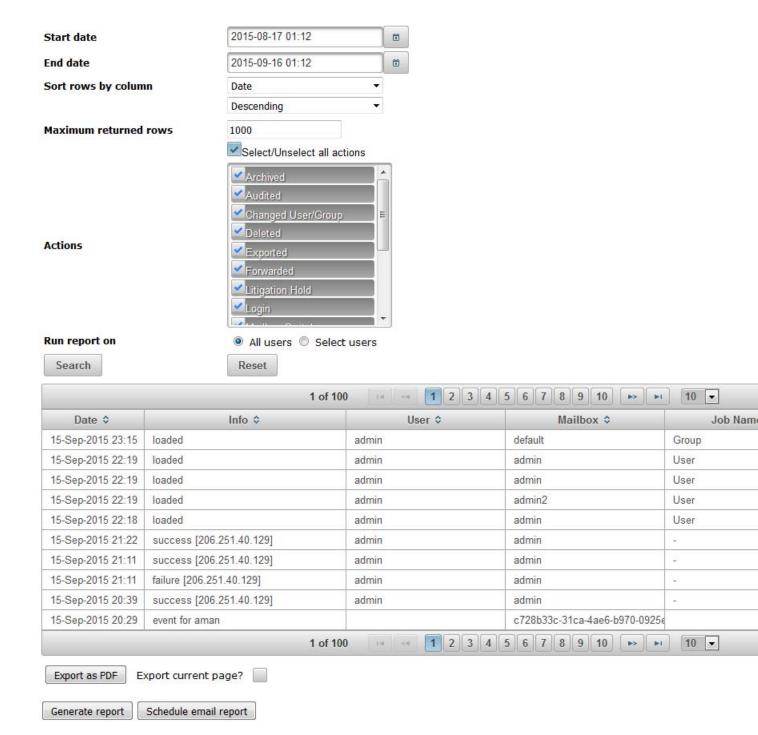
Audit Log

All pertinent actions performed by users in Retain are logged and can be discovered and printed. If Reporting and Monitoring is installed, four buttons are added to the following page, as shown below.



Audit

Report ▼ Schedule Report ▼



Accessing the audit log requires the "Access all audit logs" administrative right.

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

To view actions, select the desired actions, one or all, a date range, and a user, one or all, and then select 'Search'. The results will be displayed in a table below. All columns in this table are sortable.

The results appear below. This may be printed by selecting the 'Export as PDF' button at the bottom, and printing the resulting PDF. The PDF export button will only display what is shown in the action table, not the criteria or the interface.

If you have rights, you may click on the Action link to open the message. An auditor does not automatically gain the rights to READ messages in mailboxes that are not assigned to them.

If the Reporting and Monitoring module is installed, the 'Report' and 'Schedule Report' menu options on the top bar, and the 'Generate Report' and 'Schedule email report' buttons along the bottom will appear. These buttons link to the R&M interface to extend R&M ability to generate and schedule reports based on the selected criteria. If the Generate report or Schedule email report buttons are pressed, the criteria used in the Audit window will be automatically propagated to the Reporting interface.

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.



Overview

Documentation

Search Messages

Auditing

System Log

Reporting

Monitoring

Data Collection

Schedules

☐ Profiles

GroupWise

Exchange

Social Messaging

Mobile

Google Apps

BlackBerry

CellTrust

Workers

Jobs

Management

Users

 Device Management Deletion Management

Configuration

Server Configuration

Rules Configuration

Module Configuration

Router Configuration

Storage Manager

Stubbing Server

Reporting and Monitoring

Licensing

Logs

Clear Buffer

Logfile: /opt/beginfinite/retain/tomcat7/logs/RetainServer.2015-09-16.log

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 14:05

Sep 15, 2015 11:44:42 PM Job already completed

Sep 15, 2015 11:44:42 PM Job 44344D1D0DFD8521B8D104F36043601D will NOT be sent

Sep 15, 2015 11:44:42 PM Job already completed

Sep 15, 2015 11:44:42 PM Job 7C2C395B6BAC26D419E0713E26EB58E1 will NOT be sent

Sep 15, 2015 11:44:42 PM Job already completed

Sep 15, 2015 11:44:42 PM Job A20283E930E6E8D2858FAB40238B7CF7 will NOT be sent

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 08:00

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 16:00

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 17:00

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 10:00

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 13:05

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 14:05

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 11:00

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 11:00 Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 14:05

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 15:00

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 12:30

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 17:00

Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 13:05 Sep 15, 2015 11:44:42 PM have scheduletype: time, f_scheduleID: Daily at 15:00

About and Diagnostics Page

There is a special page in Retain containing credits for those who helped with the project. To get this page from anywhere in the administrative console of Retain, click on the GWAVA logo in the top left corner, or the 'about' link in the top right.

There are also valuable diagnostics on this page. When tuning the memory usage of the Retain components, it is valuable to refer to these diagnostics and statistics.

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



GWAVA Retain 4.0.0.0

Product Management: Jared Allen, Colin Bretagne Project Manager: Carsten Piske Architect: Dave Romanek

Lead Developers: Dave Nuttall, Ali Nikzad

Additional Developers: Christof Lorenz, Steve Orchard, Mojtaba Ebadi, Wei Han, Andro Michel, Ashwin Anathasayanan, Olga Tsibulevskaya, Swapnil Warke, Soul

Documentation: Brice Bitter

Support: Ernie Riedelbach, Ben Horne, Steve Dorrough, Ben Bishop, Chase Taylor, Daniel Cobia, Stephan Fassmann, Tay Kratzer, Dominik Niehaus, Gert Lantin Woltering, Oliver Busch, Thorsten Opitz, Perry Mason, Dean Ormsby, null

QA: Rick Liljenquist, Eric Maughan, Matt Southwick, James Hopkins, Rajwinder Kaur, Amandeep Kaur, Hakim Rasoli, Fathima Musallam, Ajay Kariappa, Neema Sat Niloofar Khoshsiyar, Suman Alungh Graphics: Tom Rains

Additional Assistance: Charles Taite, Eric Raff, Travis Doherty, Jay Parker, Preston Stephenson

Former Members: Eric Duffy, Valerie Sucher, Fei Li, Meng Yin Tao, Hadi 'Eddy' Sayar, Kevin Du, Ilya Petrov, Brad Hackett, Jason Bailey, Daron Parcell, Nischal Am Michael Bell, Art Purcell, Florian Grubler, Paul Jones, Alex 'Robert' Quintero, Manfred Bergmann, Russell Bateman, Svetlin Petrov, Joe Bertnick, Michael Sturm, Uw Joerg Buelow, Willem Bagchus, Taylor Cochrane, Raymond Hulha

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.

Diagnostic Buttons

In usual operation, none of these buttons will be used without instruction of support. The diagnostic buttons found on the diagnostics page are used in conjunction with support and can be dangerous for the system if used without understanding or instruction. The utilities contained here display or modify core settings and functions of the Retain Server. Mismanagement of these functions may render the Retain Server unmanageable and in a nonfunctioning state.

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

uptime 0d:10h:5m

currentTime Wed Sep 16 01:19:37 MDT 2015

timezone Mountain Standard Time

 retainName
 rickl-retain40

 jvmVersion
 25.31-b07

 jreVersion
 1.8.0_31-b13

javaHome /opt/beginfinite/retain/java8/jdk1.8.0_31/jre

tomcatHome /opt/beginfinite/retain/tomcat7

indexDoc 77

logHome /opt/beginfinite/retain/tomcat7/logs

osArchitecture amd64 osName Linux userName tomcat

tmpDir /var/cache/retain-tomcat7/work/Catalina/localhost/RetainServer

tomcatVersion Apache Tomcat/7.0.56

tomcatFreeMemory 2517827 tomcatMaxMemory 4415488 tomcatTotalMemory 4415488

indexEngine High Performance Index

indexHost

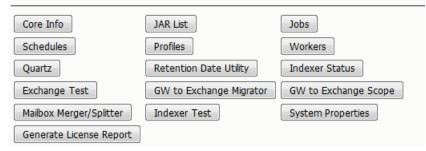
indexPort 10000 indexProtocol http

retainHome /opt/beginfinite/retain/RetainServer

retainEdition Standard
retainBuild 697
retainAvailability rc
configSchema 25
messageStoreSchema 5
indexingSchema 1
derbyVersion 10.5.3.0

indexingType High Performance Index

hibernateVersion 3.6.9.Final



The Core info, JAR List, Jobs, Schedules, Profiles, and Workers buttons initiate a download of the configuration .xml for each of those categories. All configuration tied to those items which have been configured in the Retain Server will be contained in the downloaded .xml and can be used to examine recorded configuration.

Quartz

The Quartz button displays a page containing all information and configuration for the Quartz scheduler. These settings determine when jobs start, when maintenance is, and time in the system.

Retention Date Utility

The Retention Date Utility tests the retention and purge flags in the mail systems. It is set in UNIX time stamp format, and should not be used without express guidance from Support.

Indexer Status

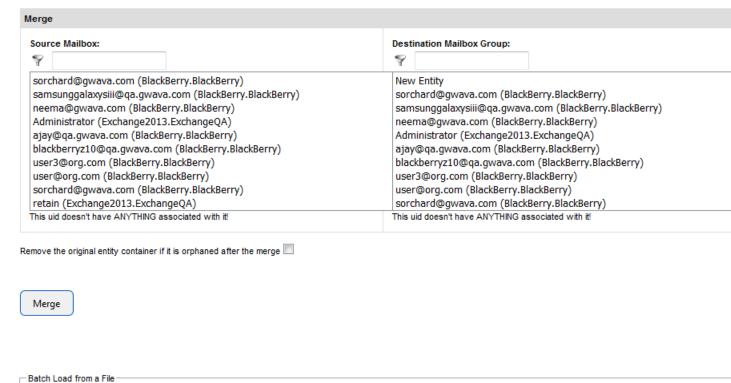
The Indexer Status displays the current health of the indexer engine along with the total unindexed items and associated memory load. There is also an option to shut off or to turn on the indexer. As the indexer should be alive at 'on' at all times during normal operation, only use this utility under instruction from Support.

Exchange Test

The **Exchange Test** utility attempts to connect to and login to the specified mailbox in the Exchange mail system, utilizing current configuration from the Exchange module, if configured. This is used to ensure that sufficient rights and access is open to Retain in the Exchange system. A target mailbox different from the Retain 'service' account must be specified.

Mailbox Merger / Splitter

This utility enables you to assign a mailbox to the virtual mailbox of another entity, it should be used with care by a knowledgeable individual. Reliable backups are recommended.



Mailboxes may be automatically merged by the Retain system if configured correctly under the Module Configuration | Mailbox Mapping options screen, but it must be done in correct order or the system will not recognize that a change has occurred and the merge will fail. Automatic merge can only be performed once.

Refer to mergesample.xml for the format.

Browse...

No file selected.

XML File:

Merge

Manually merging and splitting mailboxes is repeatable and is also the only way to undo a merge, automatic or otherwise. A batch load may be created and processed, but requires specific formatting, see the 'mergesample.xml' for format information.

Merging accounts creates a connection between two existing entities in the Retain database by merging their UID's in Retain, and allows a user with two differing email accounts, or one account repeated on two different systems, to manage their mail from within the same interface without having to search for and select a different account, or separately logging-in to two different systems.



To merge two addresses, first select the source mailbox and then select the destination mailbox. After the desired mailboxes have been selected, click the green 'merge' button.

To remove currently merged accounts, select the account you wish to sever as the source mailbox and merge it to a 'new entity'.

Because the merge utility relies on tying entities to accounts in the database, the 'new entity' simply unties the account from any other connections. Tying a user to a new entity does not create any new, floating entities or make any changes in the mail systems; it only severs the selected account from previously merged accounts.

GW to Exchange Scope

Create a new scope file or modify an existing scope file. Please select one of the two options to continue.



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

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

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

	Source	Destination	
	BenB Ben Bishop benb@qa.gwava.com	a0conn01 a0conn01 a0conn01@exchange2013.qa.gwava.com	
	BriceB Brice Bitter briceb@qa.gwava.com	a0izyu01 a0izyu01 a0izyu01@exchange2013.qa.gwava.com	
	LANG2-AR Language Arabic lang2-ar@qa.gwava.com	Administrator Administrator administrator@exchange2013.qa.gwava.com	
	User06-Projects User06-projects user06-projects@qa.gwava.com	a0john01 a0john01 a0john01@exchange2013.qa.gwava.com	
Genera	Generate Scope File		

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

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

GW Exchange Migrator



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

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

Migrating data from GroupWise to Exchange is not a straight-forward endeavor, as Exchange contains default settings which, at default, block messages from 'untrusted' source, messages with specific attachment types or attachments which are too large. (Default limit for Exchange is 10 mb.) Setting appropri-

ate trusted message types, sources, and size settings in Exchange are the responsibility of the system administrator. GWAVA Inc. is not responsible for any data that is not migrated due to incorrect settings.

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

Configure and cache the address book for both GroupWise and Exchange modules.

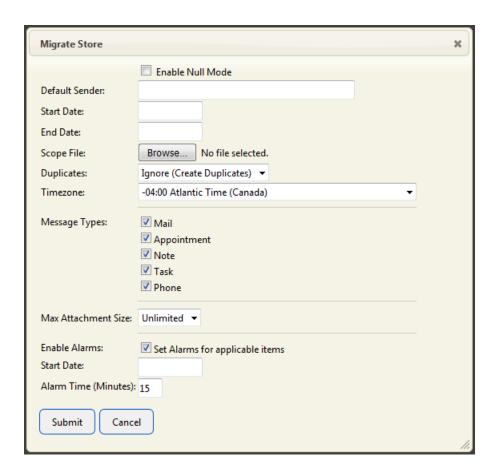
Merge the virtual Retain accounts either manually or via the automatic merge though mailbox mapping. **Immediately** run an Address Book refresh from the modules page. (If the address book is not re-cached *immediately* after an automatic merge, the merge will fail and addresses will have to be manually separated and re-connected. Forcing a failed automatic merge can be accomplished by contacting GWAVA Support.)

From the migration window, select the tools in the below order. A Scope file is required for the migration. The 'Scope Files' button links to the GW to Exchange Scope tool found on the main 'About' page, but opens the tool in a new window.



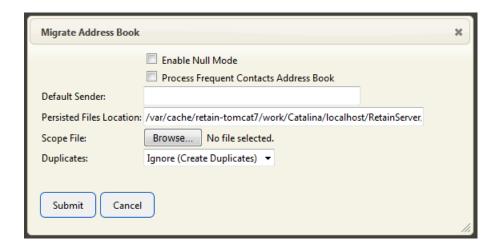
Create a migration scope profile

Open the migration tool and fill in the pertinent information. Run the migration



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

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



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

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

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

Indexer Test

The indexer test utility provides a simply and quick query to the index engine and database. This is used as a low-level diagnostic tool for Support and is best used in conjunction with Support. Normal operation does not require any use of this tool.

System Properties

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

Backing up Retain

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

- ASConfig.cfg
- License
- Index
- 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 index backup under Maintenance. Retain will automatically backup ASConfig.cfg and 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 locations to backup are:

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

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

Fine tuning GWAVA Retain

Tomcat Memory Usage Tuning Overview

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

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

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

Two Critical Issues

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

The memory allocation is shared among all web applications running under the Tomcat instance. You cannot specify a specific amount to be reserved for one specific web application. This is the major reason it is strongly recommended that Retain components (especially the Server component) run on a dedi-

cated 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:

- 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.
- Maximum Memory available to Tomcat
 - This is the most critical parameter to tune.
 - If the memory usage grows beyond the startup allocation, Tomcat will allocate additional blocks of memory in chunks as needed up to this limit. It will never return this memory to the general server memory pool. (Although memory internally will be freed and reused for Tomcat applications).
 - Tomcat guarantees that it will never exceed this memory allocation parameter. If Tomcat runs out of memory, it will try to reclaim unused memory via garbage collection. If this is insufficient, the web application will be denied the memory allocation. Unpredictable (but invariably unpleasant) results will then occur.
- Stack Size
 - For each thread (which includes each and every concurrent user request), Tomcat will allocate stack space.

This value is typically measured in KB, and defaults to 512KB.

It is allocated per concurrent users, and is in fact far too generous a number in general. 1000 users for example would take 500 MB of RAM just for stack space, before the program even allocates memory to run!

Generally we recommend reducing this number to 128k-160k. Increase conservatively, in 64 KB chunks, if you see out-of-stack-space errors in the logs.

How to change memory values

Linux Embedded Tomcat Memory Tuning

For Embedded Tomcat which ships with Retain, Edit the following configuration file: Tomcat configuration file j2ee location: <</etc/opt/beginfinite/retain/tomcat/j2ee>>

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

CATALINA_OPTS="-Xms256m -Xmx1024m -Xss160k"

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

By default, the embedded Tomcat is pre-tuned for basic functions, with 1024MB as the default for the Retain Server, and 256MB for the Retain Worker. This will not be enough in a full production environment. Please tune to fit your system needs.

Windows Embedded Tomcat memory tuning:

Go to the System Tray icon for Tomcat.

Choose configure.

Set the Minimum Memory (always used), Maximum Memory, and Stack Size.

(Retain Tomcat is pre-tuned to 1GB for the Server, and 256MB for the worker, and 160K stack size.)

Checking memory

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

UPGRADED SYSTEMS

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

Windows:

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

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?

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

A Brief Review on How Reload Works

No Helper Software Needed.

Reload runs on a Linux server. It does not use agents or helper software on the source post offices to work. In other words, no agents or TSA's are required. Reload simply connects to the server where a post office or domain is stored and then copies the data to its backup storage location.

Backups are instantly available.

Because the data is copied in its original format, the data becomes available as soon as a backup job is complete by simply running a post office agent (POA) against it (for post office backups) or a message transfer agent (MTA) (for domain backups).

Backups Have Little or no Impact on Users.

Because Reload does not use the Post Office Agent to make backups, there is very little impact on users. The POA will continue to run and service users as normal. Reload also does not use TSA software or helper agents on the live post office server. Backups can be made while the users are logged in and working.

Reload Leverages GroupWise's Architecture to Save Bandwidth.

A GroupWise post office is composed of databases and overflow files. Databases contain users' mailbox layouts and indexes and other databases contain users' authentication information. For any GroupWise item exceeding 2KB in size, such as e-mail with attachments, overflow files are stored, commonly called BLOBs (Binary Large Object).

While the contents of the databases changes almost constantly, the BLOBS are static. Therefore, in a Standard Backup, Reload grabs the databases in their entirety but only those BLOBS that have been newly created since the last backup.

Generally, the BLOBs take up almost 90% of the space occupied by the whole post office. Therefore, with a standard backup, Reload can get away with copying only 12% of the data – the databases and only those BLOBs which have recently been added. For those BLOBs that have been backed up in prior backup instances, Reload links to a master backup directory, taken the first time a backup was run, using a Linux feature called symbolic links. A symbolic link is like a Windows shortcut except that it looks, feels, and acts like the real thing.

This is how Reload can achieve VERY fast backups. In addition, using Reload to move data will save tremendously on network bandwidth compared to traditional backup systems which grab all of the data.

Backups Can be Made of Backups Allowing Centralization and Redundancy

Reload was made to backup live post offices and domains and it can also make backups of other Reload backups. The following two cases can help illustrate how useful this can be:

Consider client "A" who has two physical locations, one post office in each. This client wants redundant backups – a primary backup plus a secondary in case the primary fails.

This client installed a Reload server in each location. The servers backed up the local post office in addition to making a backup of the Reload server in the other location. Thus, each Reload box effectively had backups of both servers.

Consider client "B" who has one central data center and four branch offices. This client wants the head office to have backups of all post offices in all locations.

Branch offices 1, 2, and 3 have fast WAN links to head office but branch office 4 has a very weak connection to head office. However, branch office 4 has a strong WAN link to branch office 2.

So the client installed a Reload server in each office and one in the head office. The Reload server in the head office was set to back up the Reload servers in branch offices 1, 2, and 3. For Branch office 4, the Reload server in Branch office 2 was set up to back up the data from the Reload server in branch office 4 and then the head office was set to back this data up from the Reload server in branch office 2.

Thus, backups can make as many hops and can be backed up in as many places as you need.

Using the ability to backup one Reload server with another, you can achieve data centralization and redundancy. The redundancy also gives you the ability to use Reload for off-site disaster recovery. Additionally, for client "B", their old backup system moved all of the data every day. Using Reload, they managed to cut their network traffic by 88%.

How Retain Takes Advantage of Reload's features

Consider client "B" from the earlier example who has four branch offices and a head office. They want their Retain Server to be located at head office. So they need to centralize their data.

Without Reload, they would have Retain Workers on the branch office POA servers and the data would be sent over the WAN links. For a data collection involving "everything", all data would surely saturate the WAN links.

Plus, their backup/restore software would use the WAN links too, if they were centralizing their backups. Adding Reload to the mix, they are able to achieve huge bandwidth savings and performance gains. Reload would be set up to centralize the data to one Reload server in head office, saving immediately 88% of their bandwidth compared to their existing backup/restore system.

Next, A Retain Worker would be set up on the central Reload box to draw data from all backed up post offices.

One Retain Worker can only run one job at a time, so the post offices would be archived one at a time.

Multiple Workers on One Server

It is possible to install more than one Retain Worker on one server but this would double the hardware requirements, Tomcat memory tuning, and is limited to Linux as the platform OS. This option is built into the Linux installer and is activated by using the 'addworker' switch to the install command. (I.e. ./Retain-Install.sh addworker)

You would normally only add additional workers if you wanted to dredge more than one post office at a time.

On a Reload server, it might not be so time critical to dredge the post offices on it since there is no impact on the end users. On top of that, Reload has a special feature made especially for Retain, a special post office agent that stays up all the time, except to move to the latest backup. This way, it is always available to Retain.

So you will have to decide if it is acceptable to have the post offices dredged one at a time or if you would prefer to dredge many at a time. To do many at a time requires multiple workers.

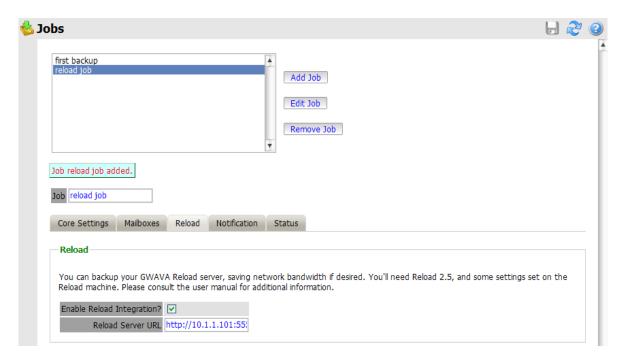
Timing

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.)

Server 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 n 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 br to reach the Worker or Server.
For administrators using the Reload interface, note the worker password assigned must be the same for each worker that contacts the Reload Server. As part setting of the Reload Server, you must know this password.
Server Protocol http ▼
Server Host Name localhost
Server Port 48080
Server Path /RetainServer
Worker Password ••••••• (Password is default)
Note: Changing any of these parameters will require recopying the worker bootstrap file!

How to Setup Reload to work with Retain

This part assumes that you have already set up and configured one or more Retain workers to collect data from your Reload box.

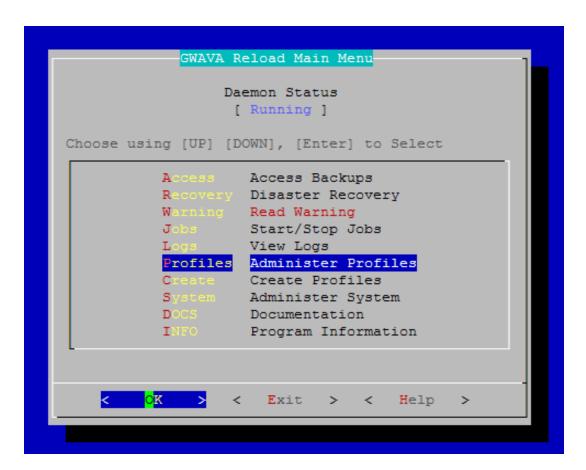
First, Reload must be set up so that the backups are available. There is a special feature in Reload for this. It calls up a post office agent that stays up all the time and it only goes down long enough to change to the most recent backup. So it will always be there with very small interruptions as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

Setting up Reload is done on a Profile-by-profile basis. Each post office that you set up for Retain to dredge from must be configured within the profile configuration menu.

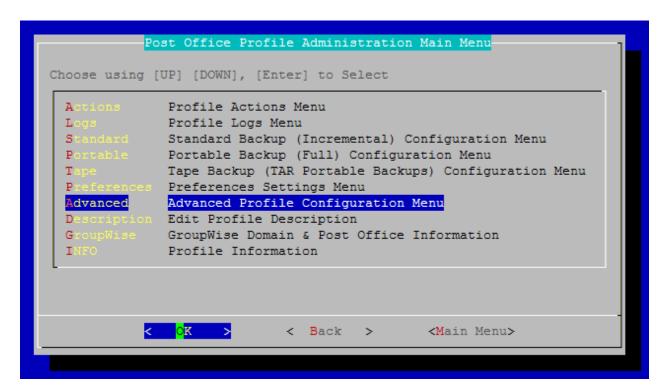
1. Start up Reload's Administration menu.



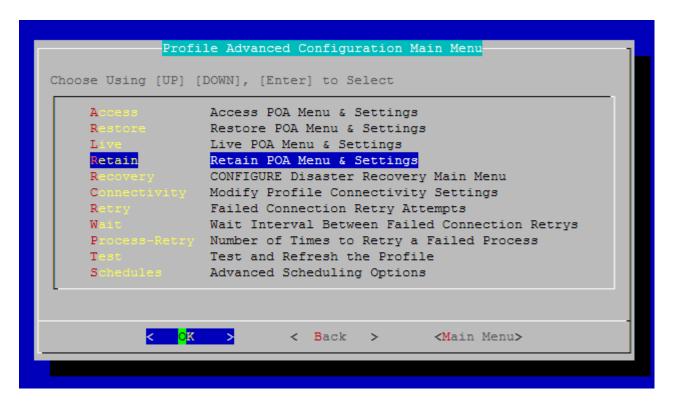
2. From the main menu, choose Profiles – Administer Profiles.

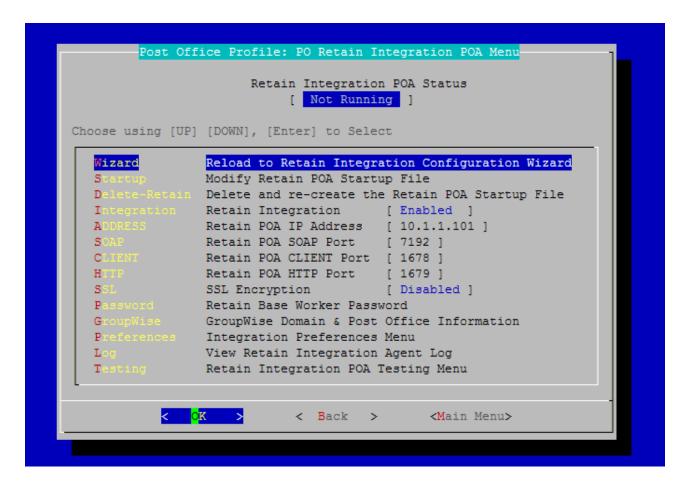


3. Choose Advanced Profile Configuration Menu



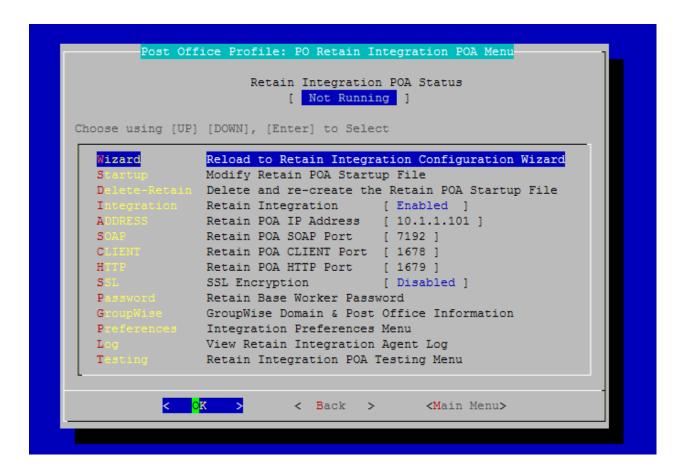
4. Choose Retain POA Menu & Settings





Now it's worthwhile examining this menu. It contains all the settings you will need to make the Retain integration work. This will be a new post office agent running and it will not interfere with the POA being used to access, backup or disaster recovery. Thus, the settings NEED to be different. The easiest way to start is to run the wizard.

Here is another shot of the configuration screen after running the wizard. You will see sample settings and the menu options will be described below.

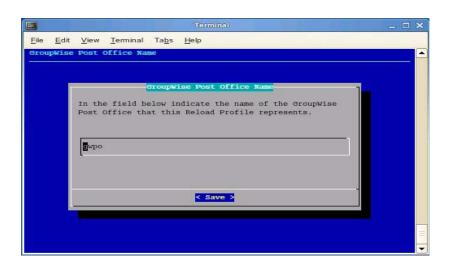


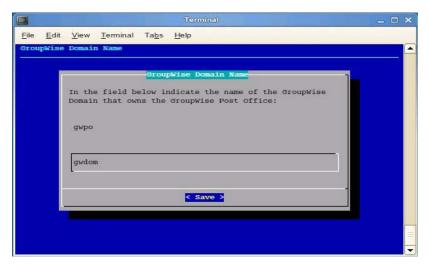
At the top, the status of the Retain Integration POA is displayed.

- Wizard: Run the configuration wizard.
- Startup:modify the startup file for the POA if you want to make specific changes to it.
- Delete-Retain:delete the startup file if you want to start fresh and configure from default.
- Integration: Enable or disable the Retain Integration
- Address:The IP address this POA will listen on.
- SOAP: The SOAP port this POA will use.
- CLIENT: The port that a GroupWise client may use to access this POA.
- HTTP: The HTTP port for this POA.
- SSL: Enable or disable SSL (Generally keep SSL Disabled)
- Key: A password Retain will use to access this POA.
- GroupWise:Specify the domain name and post office name for this POA.
- Log: View the Integration Agent Log.

The wizard will be shown below.

1. Run the Wizard





2. Enter the name of the post office and domain.



3. Choose an authentication key that Retain will use to access this POA. This must match the password you assigned to the Retain Worker.





4. Specify the IP address and SOAP port for this POA. Be sure it is unique. Some will choose one IP for the whole box with different client ports and SOAP ports for each POA. Others use the same ports but the IP addresses are different.

	Retain HTTP Port		
	In the field below indicate the HTTP port on this Reload server that the Retain Integration POA should use.		
	NOTE: Use digits, not letters		
	10080		
ı,	< OK > <cancel></cancel>		

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.

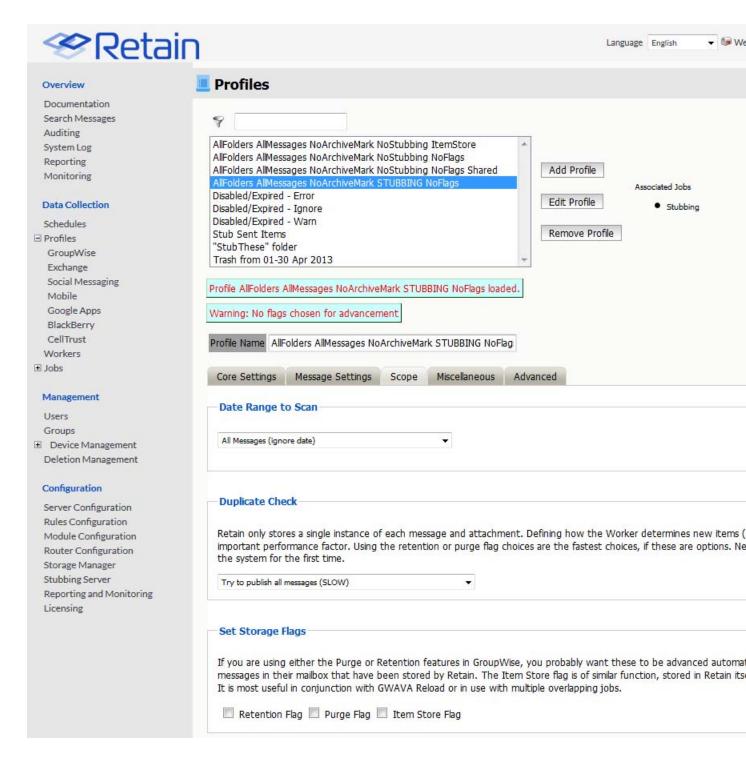
Retain will pull all necessary connection information from the Reload server. There is no need to enter these settings into the Retain Server.

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 item store flag is set in two places: Duplicate Check under the Scope tab and under Set Storage Flags under the Miscellaneous tab. The correct settings are shown.

Set Storage Flags

If you are using either the Purge or Retention features in GroupWise, you probably want these to be advanced automatically as items are stored so users may d 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 fom deleting It is most useful in conjunction with GWAVA Reload or in use with multiple overlapping jobs.

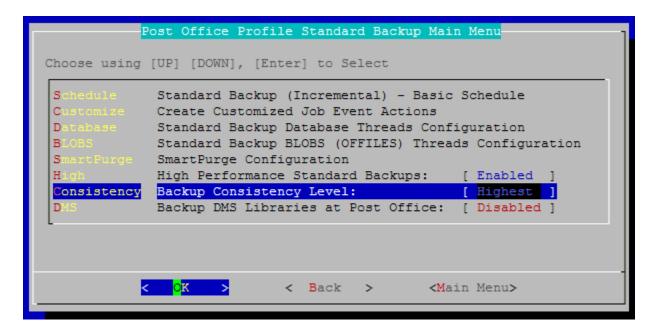
Retention Flag Purge Flag M Item Store Flag

Reload

To mitigate the chances of getting Retain Worker archive errors while working against a Reload POA, it is STRONGLY recommended that Reload is set to create highly consistent backups.

This setting is located at:

Main menu > Profiles (Administer Profiles) > Standard (Standard backup (incremental) Configuration Profile) > Consistency (Backup Consistency level): Set to highest.



This is enabled by default for new installs of Reload 2.5, but will have to be manually enabled on systems that are upgraded to Reload 2.5 You want a highly consistent backup, to make sure that you have all the blobs associated with the database. Database is picked-up first, so the blobs that are referenced in the database will be consistent with the current backup.

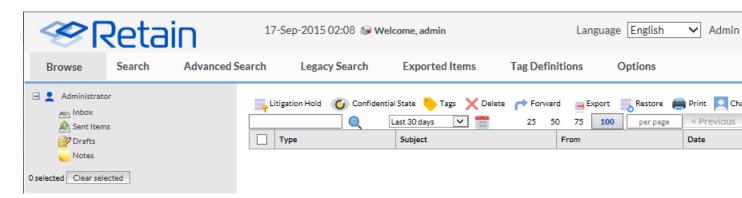
Searching Retain

The Retain search interface is the core of the information retrieval system for the Retain archive. Administrators, auditors, and users search for what they want and can view, print, export, or forward the results. General users have the right to view their own archive, and will immediately be brought to the Retain search interface, but limited to their own account.

If a user has at least one administrative right assigned, the user will see the Retain Management console on login, and must select the 'Search Messages' option from the navigation window in order to access the archive.

The full administrator interface opens with the following tabs:

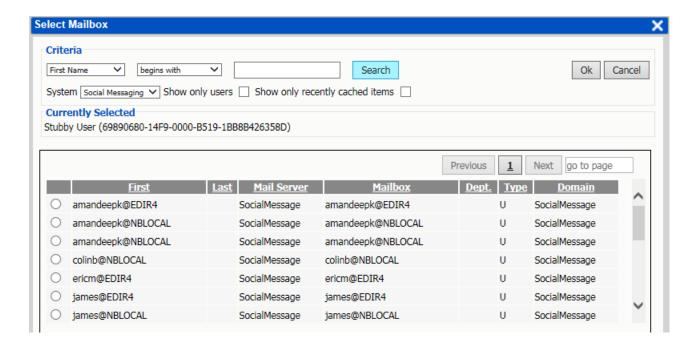
- Browse
- Search
- Advanced Search
- Legacy Search
- Exported Items
- Tag Definitions
- Options



Browse

The Browse tab immediately grants access to the selected mailbox. Individual users will only see their own mailbox, while users with the administrator right to search all mailboxes will have the option to change mailboxes to another user.

To change to a different mailbox, select the 'Change Mailbox' button in archive toolbar. Once clicked, the 'Select Mailbox' dialog will open.



The Select Mailbox tool asks for specific information. The Select which mailbox you want to see by clicking on its radio button and clicking OK.

When searching for a mailbox, the system of the mailbox must be selected as mailboxes from different systems may have the same user name or criteria. Select which mail system the desired user belongs to, specify any further criteria, or leave the criteria blank to display all possible mailboxes from that system.

(If Retain for Social Messaging is set to anonymous user, all Social Messaging data will be contained under the single user '?@?', and separate user names and pertinent information is contained in the 'from' dialog. Otherwise, individual user accounts will be displayed.)

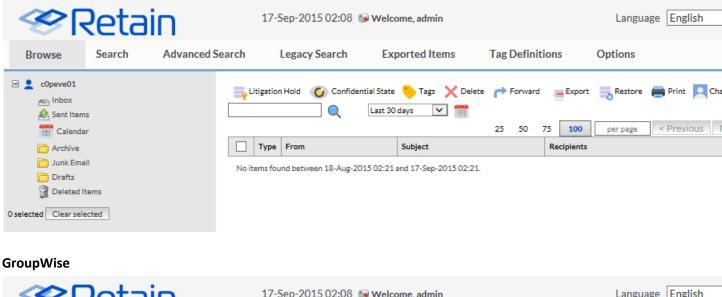
If the search results are extensive, the system will have a 'Next' or 'Previous' button at the bottom of the search results window, which displays the next set of results.

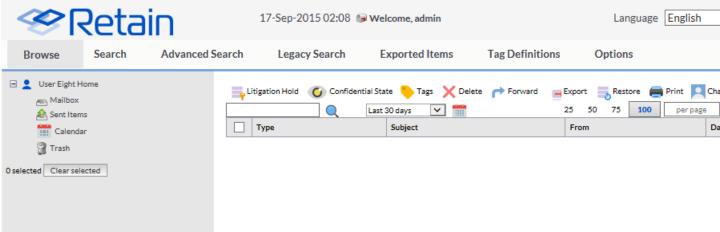
Refine search parameters to reach a manageable search result.

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

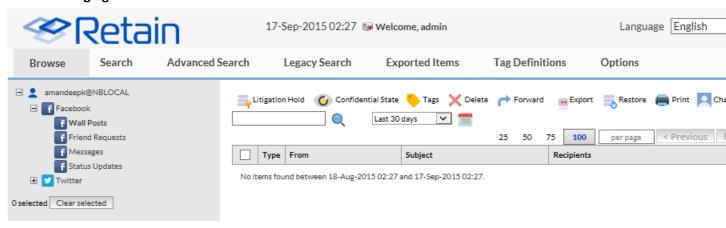
After selecting a mailbox, click 'Ok' to load that mailbox into the viewer.

Exchange





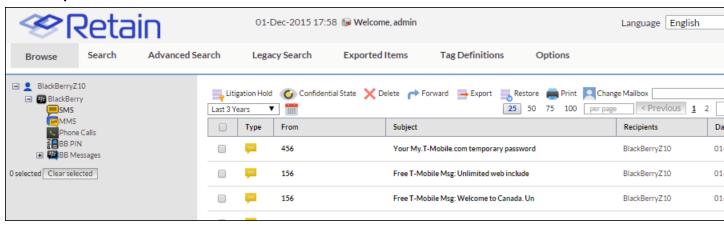
Social Messaging



Google Apps



BlackBerry



CellTrust



Toolbar



The toolbar has many different tools. Which tools are visible to each user will depend on what user rights they have been granted.

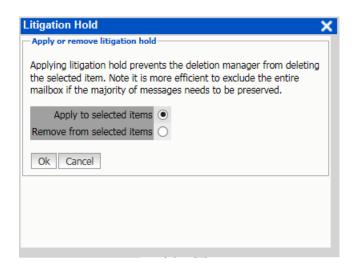
Date range Last 30 days 💟 🛗

Pay special attention to the date range tool. Only messages which fall into this specified time frame wil be displayed. The date range can be any of the drop-down options or it can be a custom range as selected through the 'custom' calendar button.



Litigation hold ____Litigation Hold

The Litigation hold button applies the litigation hold property to messages and items. Items with a litigation hold cannot be removed from the interface. To apply or remove the litigation hold from items in the archive, a user must have the litigation hold right.



To apply or remove a litigation hold, first select the desired item or items, then click the 'Litigation Hold' button to open the ligitation hold window.

Select whether to apply or remove the hold. Select 'Ok' to continue. Items which have a hold placed on them display a miniature hold icon next to the message icon.



Toggles the confidential state. Confidential items are only visible with Admin rights, or to those specified by the user under the Options | Confidential Exceptions tab. Items which a user has marked in their box as Confidential may be viewed by other users who have the 'See Confidential Items [other mailboxes]' admin right.

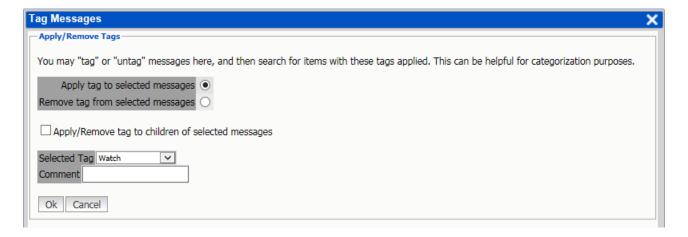


To apply a confidential state to an item, select the item then select the 'Confidential State' button. Items which have been marked confidential display the confidential icon below the message icon.



Tags 🐎 Tags

The tags icon allows the application or removal of tags from items. This is only visible with rights. Tags are an informative note which can be attached to any data item in the search messages interface. To apply a tag, select an item and click the 'Tags' icon.



Select the tag which is to be applied. All tags open to use by the user will be displayed. There is no limit to the amount of tags an item may have. An additional comment may be added to the tag for the specific item or items selected. Additionally, any associated items or messages, 'children', such as all messages in a thread, may also be selected for tag creation. Tags must be created previous to applying them to messages. See the Tag Definitions tab for more info.

Messages with a tag applied to them, either a global or personal, will show a miniature tag icon next to the message or item icon.

Delete X Delete

The delete option allows users to mark items for removal from personal archives. This right is only visible with rights. To delete an item, highlight the item and then select the 'delete' button.

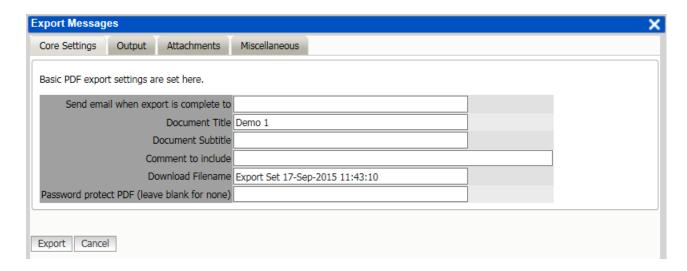
The Delete right is not a default right, it must be specifically added to users or groups. Messages or data marked to be deleted will not necessarily be immediately removed, but will be queued for deletion. The actual item removal will depend on server availability and the amount of items queued for deletion.

The Forward option allows users to forward messages out of their archive to a specified address. This right is only visible to users who have the Forward Messages right. Select a message or messages and then the 'Forward' button. Specify the address or addresses the messages should be forwarded to, a desired subject, and any comments. Finally, select the 'Forward Messages' button and Retain will send the message via SMTP.

Export Export

The Export option opens a PDF export window to create a PDF of the selected messages and emails a notification to the specified email address. Messages will be displayed as navigable files in 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.

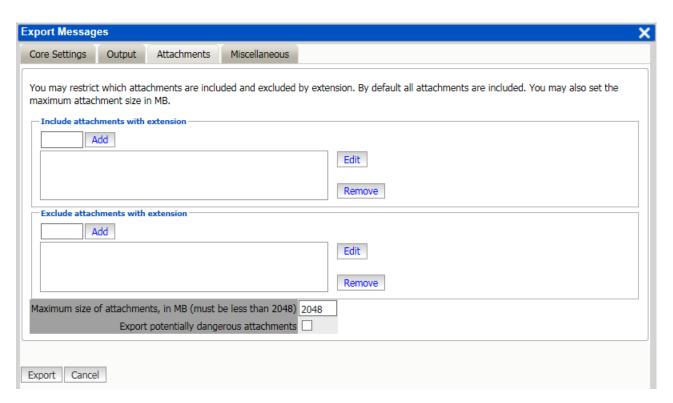
Select the desired message or messages and then the 'Export' button.



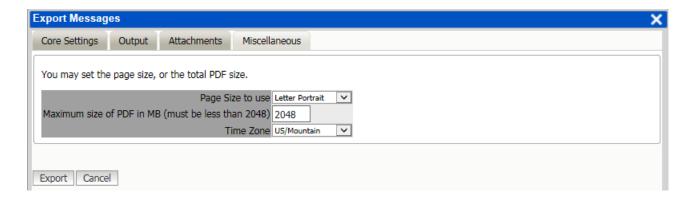
Fill out the various options as desired.



File output and included fields can be edited



Select which attachments to include and exclude from the PDF.



Additionally, any restrictions to the size or time zone for the PDF. If the archive size does not exceed the maximum size, then the archive will be contained in one file. If the archive is larger than the defined limit, it will be exported in multiple PDFs. Time zone information will be used for formatting any time and date information in the email headers. The time written to the cover sheet and date column in the PDF table of contents is written according to the time zone of the Retain Server.

When complete, select 'Export'. Depending on the size of the selected archive, the export may take some time. The Export will notify by email when it is complete if a notification address is provided. Currently running export jobs are also listed under the 'Exported Items' tab. Once complete, the PDF is available under the 'Exported Items' tab. Locate the desired export and select the disk icon to download the finished PDF.

Restore Restore

This option restored messages to the message system. To use, select the desired messages and select the 'restore' button. Once there, select the target message system and mailbox. (If the user has rights to restore to any mailbox, the target mailbox must be selected. Select the 'Restore Message' button and messages will be restored to that account.

Restoring a message which currently exists in the target mailbox results in a duplicate message existing in that mailbox. It is recommended to only restore messages that have been deleted from the mail system, but exist in the Retain archive.

Print Print

This option allows users the ability to print selected messages. To use this option, select the desired messages and then select the 'print' button. The 'print' function utilizes the user's local print configuration; if there is no printer accessible to the local workstation, print will not function.

Change Mailbox Change Mailbox

This option allows users who have the administrator right to see all mailboxes, or who have proxy rights to another mailbox, the ability to switch mailboxes.

Quick Search Q

The Quick Search is a way to rapidly search for data in the subject field of messages in the selected mail-box.

Results are automatically set to 100 per page. The selector allows the user to specify how many messages are to be displayed on each page. There are predefined values, or the user may put their own dictated limit in the field provided.

Previous/Next < Previous 1 2 3 ... 5 Next >

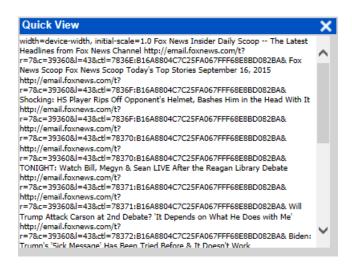
The Previous/Next option allows users to navigate through the search or browse result pages. The total number of pages is displayed next to the 'Next' button. Select to immediately move to the desired page

Go to Page go to page

The Go To Page option allows users to specify what page they want to jump to. select the option box, specify the page number desired, and hit 'enter' to immediately go to that page.

Quick View

The Search interface incorporates a quick view option. This option allows users to quickly view the body of a message without having to actually open the message in the interface. To access the quick view, click on the message or data item icon next to the check box.



Once selected, the quick view window opens to quickly display the message contents.

The Quick View is a display only function and there are no actions available from this interface. When complete, close the window by clicking on the 'X' at the top right corner of the window.

Selection and sorting



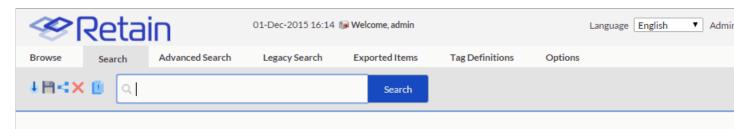
The top column bar contains all the different column headers and is located directly beneath the toolbar. This tool contains a select all/none checkbox as well as sortable column headers. Every column in this bar may be sorted by order except the 'Mailbox' and 'Folder' columns. (They are multi-value attributes and thus are not sortable.)

Browse Tab Defaults

By default, the browse window remembers which mailbox last browsed was, and will return to that mailbox when opened. When Administrators first login, they will be taken immediately to their mailbox. Afterwards, the browse will remember their last location.

Search

The search in Retain 4 is designed to work much like a Google search, with anticipated auto-correct search terms. The auto complete suggestion feature requires at least 3 characters to be specified before it will begin to work. The Search tab will allow searching across the entire Archive for administrators, while users will only be able to search their own archive.

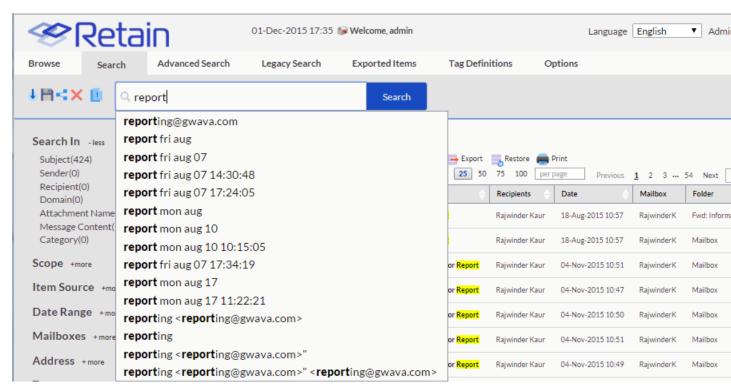


The Search tab allows for management of searches. This is accomplished by Saving searches for later use. These searches can then be shared with others, loaded later, removed when no longer desired, and trimmed to remove duplicate data.

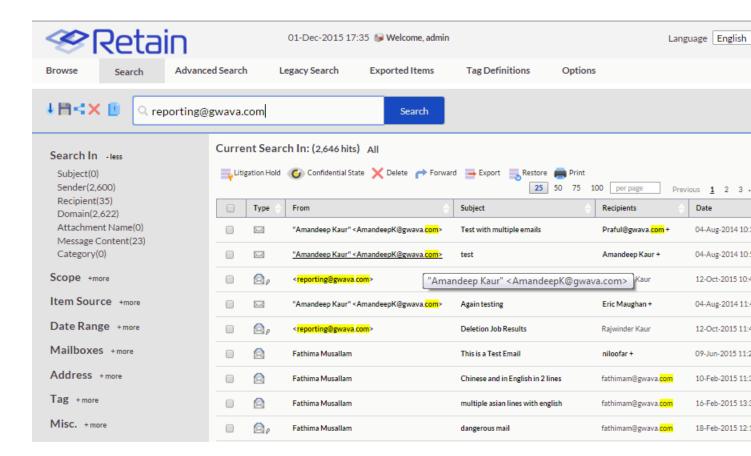


The icons for these functions are on the right, next to the search field. In order: Load, Save, Share, Delete, and collapse duplicates. Collapse duplicates removes all duplicate results but only functions with Group-Wise mailboxes.

To begin a search, type the desired search term into the search query field, or load a saved search. Results will populate before a query has been selected.



Select a final query, either by hitting enter, clicking with the mouse, or using arrows and hitting enter. Search results will display the results with the search term highlighted for each message. The message type, Sender, Subject, Recipients, date, mailbox, and folder are all displayed.



Once search results have begun to populate, the left hand scope pane is populated with limiting and filtering options. The side bar also has numbers next to each section, indicating how many hits there are for each particular topic. The hits are total numbers of matching instances, not items. So if a message states a search term several times in the message body, it will be counted as that many hits even though it is only one message.



The Search In criteria limits the area in the message or data where the search is performed. The Subject indicates hits in the Subject field. The Sender field contains the sender of the message or data item. The Recipient is the recipient of the item or message. The Message Content will only search the body of a message or item for the search term.

Searching exclusively for the domain will be effective with search terms if a complete domain is provided, otherwise the term is recognized as text. If the top level domain is not known, (.com, .org, .edu, .etc) then

the search term should use an asterisk afterwards. IE. Searching for gwava.com will yield good results, as will searching for gwava.* or gwava*, though results will vary.

The Scope criteria option limits the type of message which is to be searched. All item types are available in the scope term. Again, the number of available hits is displayed to the side.

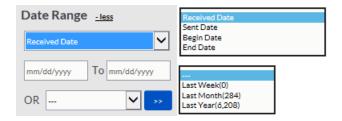
Scope -less
Mail(19,240)
Phone Message(0)
Appointment(0)
Task(0)
Note(0)
SMS(0)
MMS(0)
Phone Call(0)
BB PIN(0)
BBM(0)
Wall Posts(0)
Requests(0)
Chats(0)
Likes(0)
Events(0)
Status Updates(0)
Comments(0)
Searches(0)
Group Posts(0)
Web Mail(0)

The Item Source criteria option limits the results to a particular source. This source can be limited to show 'received', 'personal', 'sent', or 'draft' items.

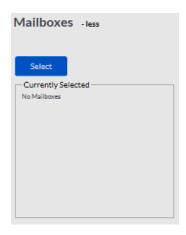
Item Source -less
Received(19,214)
Personal(0)
Sent(8)
Draft(0)

The Date Range criteria limits the time frame of item's creation. Only messages which conform to the date range selected will be displayed in the results field.

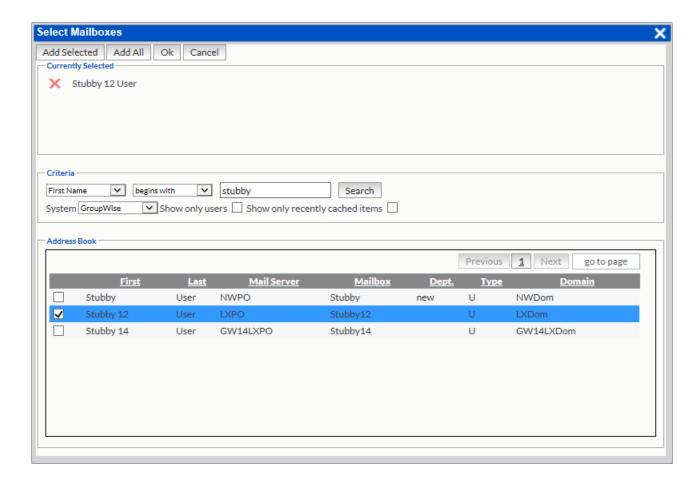
The date range may be specified for any of the dates corresponding with an item. The range may also be selected form the drop-down menu, with preconfigured times for last week, month, or year. These time frames are for the past 7 days, the past 30 days, or the past year, not the previous calendar time frame.



The Mailboxes criteria limit which mailbox or mailboxes the search will pull results from. To add users to the selected mailboxes list, click on the 'Select' button to launch the mailbox selection window.



The Select Mailboxes window allows for searching of every mailbox in the system to limit the search. Mailboxes must be searched for by system and specified criteria. The results of the search are displayed below, while the active selected mailboxes can be added to the dialog through the use of the 'Add Selected' button along the top. Alternately, if the 'Add All' button is clicked, it will add all mailboxes displayed in the search results, the Address Book field. Addresses which have been added to the top field may be removed by selecting the red 'X'.



Once the desired addresses have been added to the top window, select the 'OK' button to load them into the search pane.



The Address option limits the results to a selected address. The addresses available are displayed below, and may be selected Addresses in the window are dictated by what is in the result set. Selecting an address adds it as a filter to the top of the search window, or the user may specify an address manually. Multiple addresses may be added to the search window at a time. To remove an address filter, select the 'X' next to the active address, and the result set will be reset.



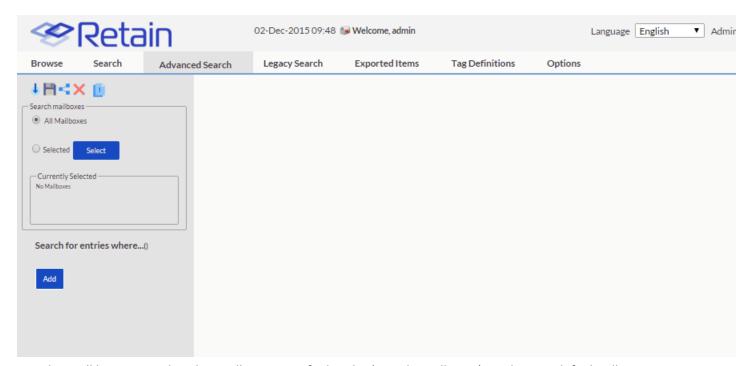
The Tag option limits the search to items which have been tagged with a specific tag definition. The tag definitions may be personal or global. Tags must be specified in advance and applied before this option will work. Select the desired tag to limit the result set.



The Misc option contains limits for the Litigation hold and the Confidential tags. They have two settings: True or False. A Setting of 'True' restricts all results to only items which have the selected tag.

Advanced Search

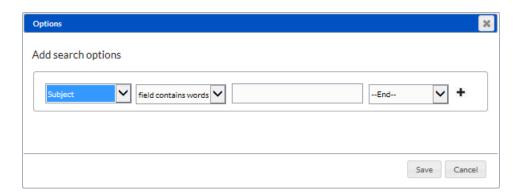
The Advanced Search tab contains the ability to specify vast amounts of criteria and combine search terms to exclude and include various searchable items to retrieve specific information. This search works better, the more you know about what you are looking for, as it allows fine tuning of criteria.

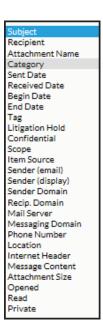


Searches will be restricted to the mailboxes specified in the 'Search Mailboxes' window. By default, all mailboxes are set for searching. To limit the search to the mailbox or mailboxes specified, click on the

'Select' button to open the 'Select Mailboxes' window. The 'Select Mailboxes' window functions exactly the same in advanced search as it does in the standard search.

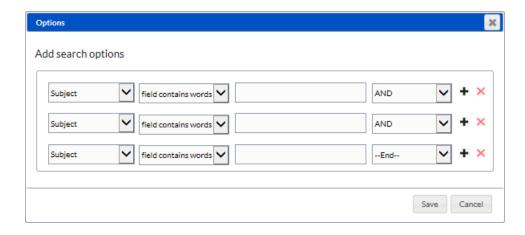
The system will begin to display results as soon as either the search mailboxes have been specified, or new search criteria has been added. To add new criteria, select the 'Add' button.





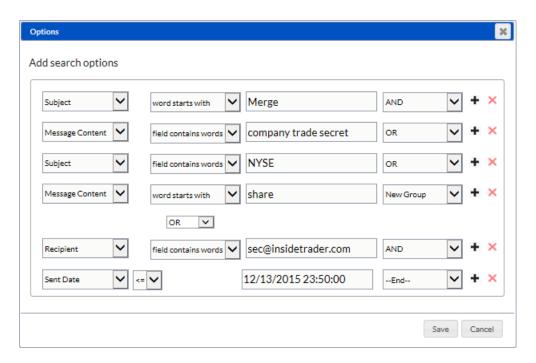
The Search Criteria contains the ability to specify where to search, operating criteria - (word ends with, word starts with, field contains words), and the desired search terms. The list of search items and fields available to be specified in the drop-down list is shown. Each variable on the list is tied to appropriate search operators, (date range allows the specification of a date, Confidential tags have a true/false operator, etc.)

In addition, the interface allows for no limit of search terms. additional terms may be added to he search criteria and connected to the previous search terms. Additional criteria may be logically connected with 'and', 'or', or 'new group'. To add a new search term and criteria, select the '+' directly to the right of the existing search criteria.



By default, when a new search term is added, it is automatically 'and-ed' together with the previous search term. This allows users to be able to build complex search terms to fit known data.

When building complex search criteria, it is critical to know what you are looking for. For instance, if an insider trading tip was suspected, and the recipient was known as well as some details about the message and when it must have been sent by, the following search could be compiled:



In this search, any message sent which stated 'merge', or 'merger' in the subject, and contained a known company secret in the message body, or, discussed the name of an executive involved, would be displayed. In addition, the search would also grab any messages sent to the suspected contact before the merger date. Additional criteria which could be added includes the company's stock listing or any further details pertaining to the proposed leak.

Search for entries where(18 hits)					
Subject word starts with	Merge	AND			
Message Content field contains words	company trade secret	OR			
Subject field contains words	NYSE	OR			
Message Content word starts with	share				
OR					
Recipient field contains words	sec@insidetrader.com	AND			
Sent Date <=	12/13/2015 23:50:00				
Edit All Delete All					

To begin the search, select the 'Save' button at the bottom of the query window to perform the search. The active criteria is now listed in the left pane, and may be edited or removed. To add criteria, select the 'Edit All' button to add to or refine the search criteria.

RegEx and Wildcards

Both the Search and Advanced Search contain limited support for Regular Expression searches. To use Regular Expressions, simply put the desired regex string into the criteria window, denoted by a '/' on either side of the regex. if the '/' is not used, Search will not recognize it as regex.

The Supported Regex characters are:

Syntax	Description	Example
,	Zero or one occurrence	/abc?d/ matches abd, abcd
*	zero or more occurrences	/abc*d/ matches abd, abcd, abccd
+	one or more occurances	/abc+d/ matches abcd, abccd
{n}	n occurrances	/abc{2}/ matches abcc
{n,}	n or more occurrances	/abc{2,}/ matches abccc,
{n,m}	n to m occurrences, including both	/abc{2,3}/ matches abcc, abccc
[characters]	character class	/[abcde]{2}/ matches ab, ac, be
[^characters}	negated character class	/[^abcde]{2}/ matches fg, hi, jk
-	character range, including end- points	/[a-z]{2}/ matches ab, ac, bz
	any single character	/ab.k/ matches abck, abdk, abek
	union	/(ab cd)/ matches ab, cd
&	intersection	/([a-z}* cd+)/ matches cd, cdd, cddd
@	any string	/@/ matches abc is abc
<n-m></n-m>	numerical interval	/<99-102> some dr/ matches 99 some dr, 101 some d

Accepted wildcard search symbols are: *, "", and ?. The '*' denotes 'any character or characters', while the '?' denotes 'any one unknown character'. The double quotes("") denote an exact phrase, and only

that exact phrase. For instance, a search for the term "We all love spoons" will fetch a result of that exact phrase, however it will miss a sentence stating 'We all love spoons for our ice cream', or 'We don't like sporks but we all love spoons.' Adding asterisks inside the double quotes will allow both omitted sentences to be retrieved: "*we all love spoons*". While the '?' works like the asterisk, it is only for a single unknown character. This is particularly helpful when searching for exact phrases where the terms may be misspelled. Ie. "Their going" would miss a misspelled 'there going'. However, if the search term were "The?? going", it would catch both.

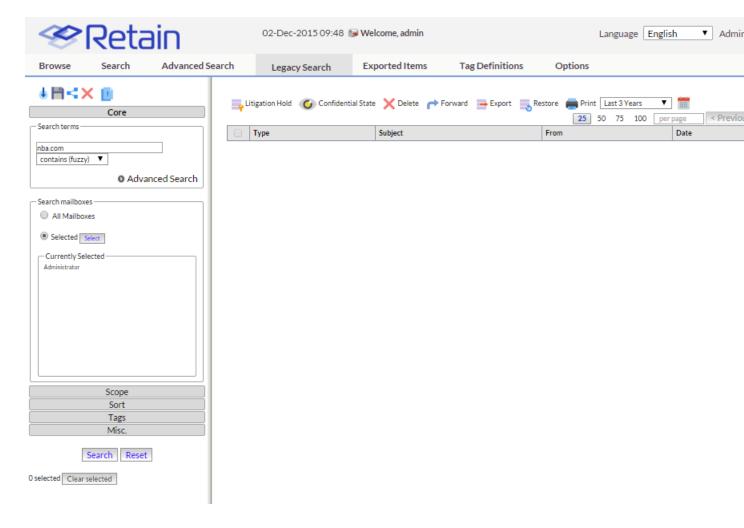
Special Characters

The Search has a list of special characters which cannot be searched for, and will cause erratic results with search criteria. The list of non-supported characters is: $@,+,|,[],{},(),",\downarrow,\#,\&,\sim$. All of these characters are viewed as delimiters, and will break up the query. They are not supported.

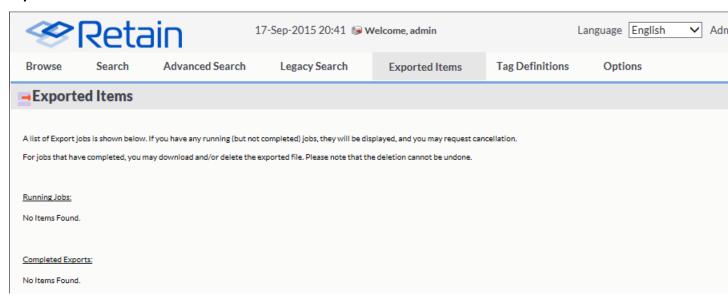
As a result, if a special character is to be searched for, it must be placed in double quotes. Ie. searching for bob@gwava.com will result in any user at the 'gmail.com' domain. To search for bob@gwava.com, the exact phrase must be specified; ie. "bob@gwava.com"

Legacy Search

The Legacy Search tab provides access to the older search engine, and is included as a courtesy for those who wish to continue using it. While the UI has had an update, the functions and search process remains the same.



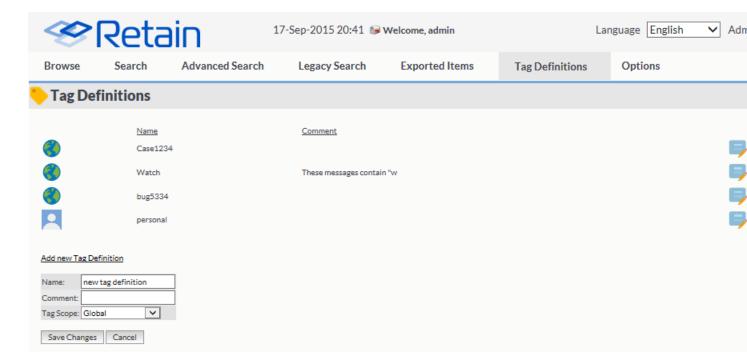
Exported Items



The Exported Items tab shows the export jobs which are currently running, or have been run in the past, on the system. If an export job has completed, this tab will contain the file and provides a link for download of the completed export job.

Tag Definitions

The Tag Definitions tab allows the creation and removal of Tags, their automatic comment, and name. Tags are an informative note which can be attached to any data item in the search messages interface. There is no limit to how many tags any one item may have applied to it, and there is no limit to how many tags a user may create. In addition, tags are also a serachable item, making this one of the most versatile ways to add long-term identification for items in the data store.

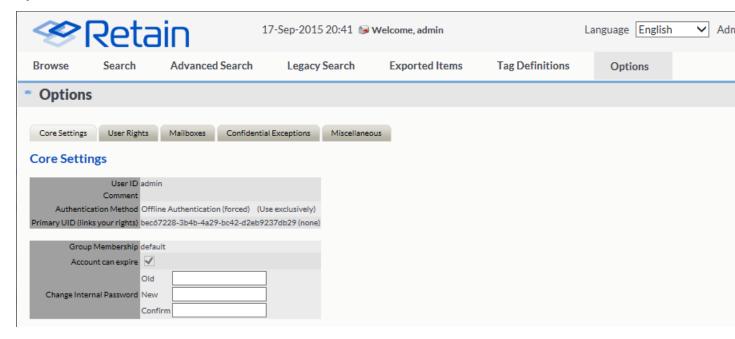


Before the tag icon will appear on the in the search interface tool bar, there must be at least one tag defined. To define a tag, enter the tag name and initial comment if desired, then, if the user has permissions to do so, define whether the tag is personal or global. Once saved the tag is available for use.

Global tags are tags that any user with the rights to see global tags will be able to view and apply. Personal tags are limited to the user who created them. Only tags visible to users will be available to be searched for by that user.

Any tags created or subject to manipulation by the user logged-in will be displayed under this tab. To apply a tag to a message or data item in the search messages interface, simply select the data item or items, and then click the 'Add / Remove' tag button in the tool bar.

Options



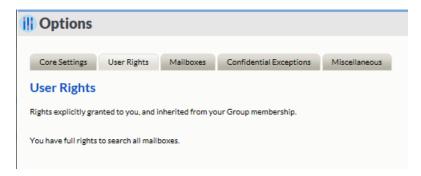
The options section here is exactly like the section in the Administration | Users section. These settings here are specific to the currently logged in user. The current loaded user is shown next to "User ID" located below the Core Settings tab.

Core Settings

Among the information displayed is:

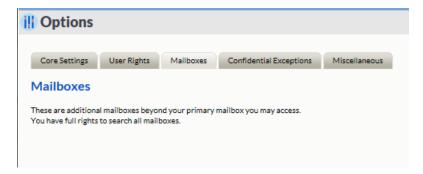
- Your User ID (internally generated by Retain)
- Your Group Membership (you cannot change that here)
- Your authentication method (you may change this if you have rights to)
- Your offline password if you don't use SOAP to authenticate
- You may change the password only if you have rights to change it

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.

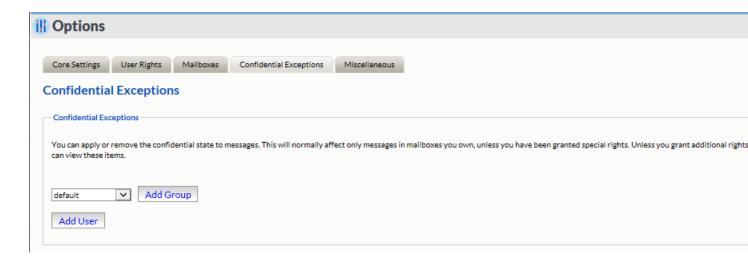
Mailboxes Tab



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.

Confidential Exceptions

When a user marks a message as 'Confidential' in the archive, the message becomes invisible to all except administrators who have been given the confidential right, or any user which is specified in the exception list.



The 'Confidential Exceptions' allows users to add any necessary exceptions to the confidential tagging. Confidential tagging may be applied to protect sensitive. However, sometimes this information may need to be viewed by others and instead of granting that user rights to see confidential items for all users, a user may apply that right to only their items.

A group or individual user may be added to or removed from the list.

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
Date/Time Fo

Date/Time Format Display Number Message Age Display

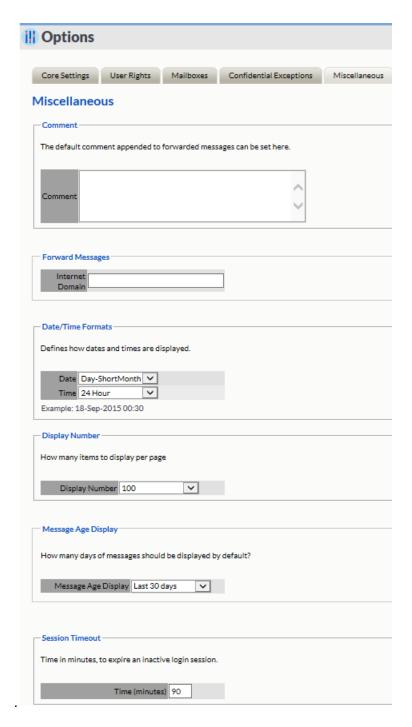
Session Timeout

Default comment for fowarding messages

How you want the dates and times to be displayed

How many items to display per page

Default date filter for searching - may be changed 'on-the-fly' Inactive session timeout - set between 5 and 60 minutes



Outlook Plugin

Overview

The Outlook Plugin integrated the Retain storage directly into the Outlook client. Users will have a new tab and ribbon added to the Outlook client offering connection to their Retain storage. Messages may be retrieved and added to the user's local machine and administered through the Outlook client. Message data may be searched in both the Retain storage, and the locally restored messages. The client will not modify Retain administration nor will anything done through the Outlook Plugin modify the Retain Server or the Retain Storage. The Outlook Plugin is provided for both individual and distribution installations.

Supported Software

Outlook 2010, 2013 32-Bit and 64-Bit

NOTE: Whether the 64-Bit or 32-Bit version should be used is determined by the bit version of Outlook installed, not the version of the OS. To check the version of Outlook being used;

2010: The information on the system is located under File | Help | About Microsoft Office

2013: The information on the system is located under File | Office Account | About Outlook.

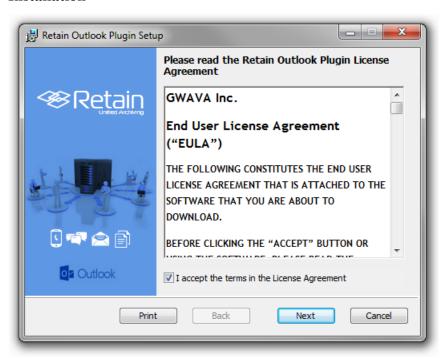
System Requirements

Visual Studio Tools for Office v. 4

.Net 4 or higher

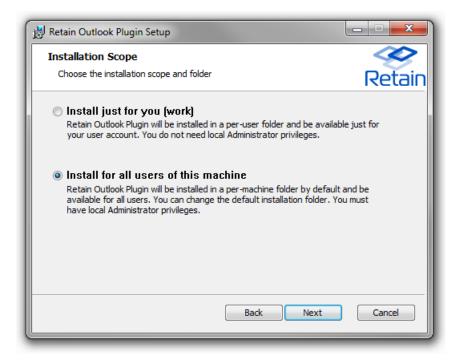
NOTE: The .EXE installer bundles the prerequisites and can be used to install both prerequisites and the plugin at the same time.

Installation



Installation of the Retain Outlook Plugin is straight forward. If unsure that prerequisites have been satisfied and installing on individual workstations, use the .EXE bundled installer.

If deploying from a distribution server, utilize the .MSI installer package. In addition, silent install commands are provided. A full list of silent install commands can be found in the appendix. Make sure to utilize the correct version of the installer to match the version of Outlook: 32-Bit or 64-Bit.

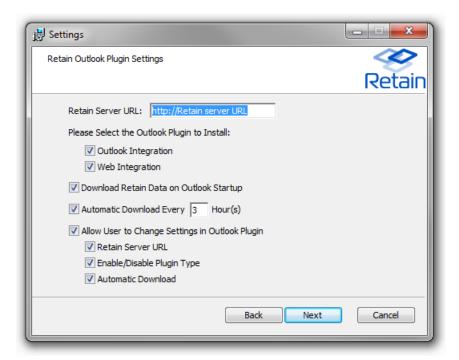


Begin the installation, accept the license agreement, and click 'Next'.

Select whether the installation will be system wide or restricted.

The Retain Outlook Plugin may be customized during installation, to only allow certain functions once installed.

Retain server URL: this is the connection address for the Retain Server. Do not put the '/RetainServer/' suffix on the end of the URL. Use the following syntax: http(s)://(Retain Server URL)

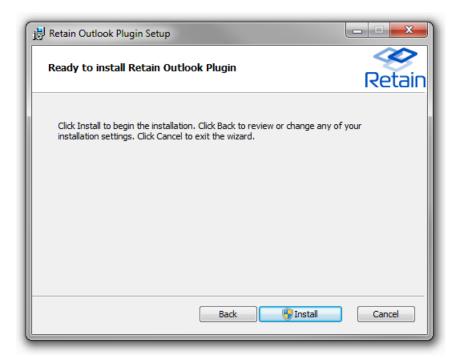


Select Outlook Plugin: There are two options for the plugin. The Web integration is the first plugin version, which links to the Retain Server Web Interface. The Outlook Integration places the Retain tab and ribbon in the Outlook Plugin and the rest of the options apply to the Outlook integration.

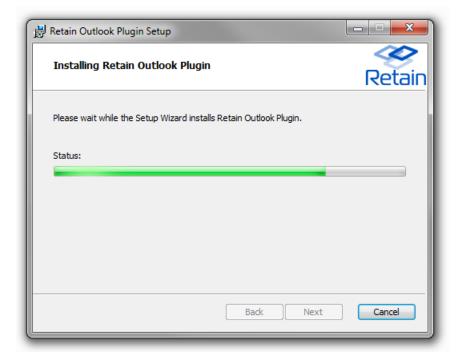
Download on startup: This tells the Retain plugin to match the local data with what is in the Retain Server. This does not push any data from the plugin to the Server. The plugin may be configured to download at start, or on an automatic time period. The setting is designated in hours.

Allow users to change settings: If checked, the following selected options are available to be changed by the user in the plugin. If this option is not checked, the options will be grayed-out in the settings section of the plugin, and cannot be modified.

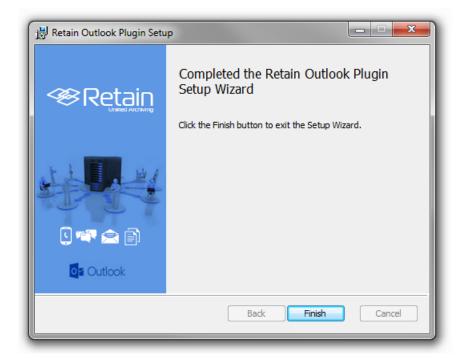
Configure as desired and select 'Next'.



If any settings need to be changed, use the 'back' button and change then. Select the 'Install' button to begin the plugin installation. (Installation will require admin rights.) If using the bundled installer, the installation will check for prerequisites and install them if they are missing.



Wait for the plugin to install.



Once the install has completed, select 'Finish' to exit the installer.

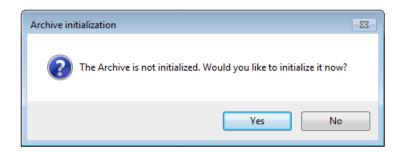
Features and Use

To access the Retain Outlook plugin, open the Outlook client.

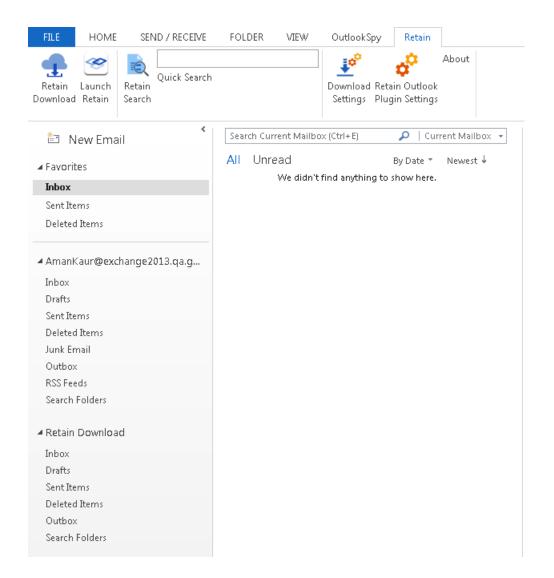
In the Outlook client, a new Retain tab is available. The Retain tab provides access to the plugin and Retain archive settings and features.

Initialize the archive.

Users now can search their Retain archive, import messages from Retain to their local system, and manipulate those local messages with the Outlook plugin.



The first time that the plugin is initialized you, it will ask to initialize the archive. Initializing the archive creates a new folder tree in the tree navigation pane on the left. The folders here are the default folder set for Outlook. Any new folders created after download, are there as they exist in the Retain Archive. Once installed and initialized, the Retain tab, ribbon, and folder structure are available.



To access the folder structure, simply select it as you would the normal folder structure in Outlook.

Download Settings



The Download Settings contains the criteria which dictates what message data will be added into the local storage and made available in the Outlook client.

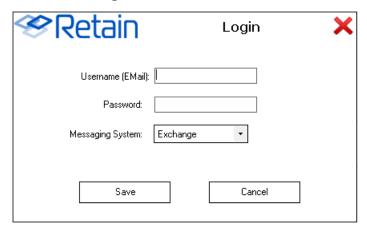
The download time limit dictates the time frame of messages which will be displayed in the Outlook client.

Messages may be limited by total number, as may the search results. (Search results limit the number of results by default to only display 500 messages.)

In addition, the type of data may be limited to only the types selected, or all types.

What data is added to local storage can be limited by age, number of messages, or item type. 'Save' saves your settings. However, before the download can be performed, the login information must be provided. Select the 'Account Settings' button.

Account Settings

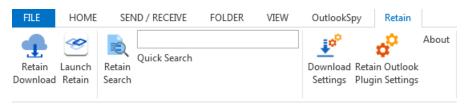


The account settings are simply the login credentials of the account which is to be used for the download. Input the email and password, and specify the message system to be used. GroupWise usernames must not be the full address. If inputting an Exchange username, enter the entire email address. Save settings.

Clear Local Archive

The Clear local archive button does exactly as it sounds, it will empty the local archive, Retain, folder tree of all message data items. Any folders which were part of the Retain archive, and not part of the default folder set of Outlook will also be removed. The default set of folders and the Retain folder tree will still be present. To completely remove the Retain outlook plugin, please uninstall the plugin from the system's Control Panel.

Retain Download



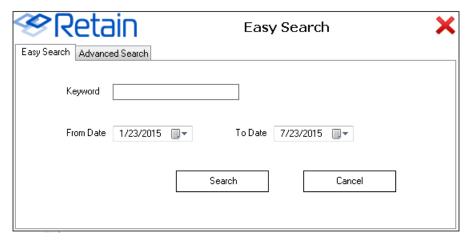
The Retain Download button initiates a download of the users Retain archive according to the download settings. Download settings and the user credentials must be configured before the download will complete.

Retain Search

Searching the Retain Archive with the Retain Plugin is simple. Select the 'Retain Search' button from the Retain ribbon.

The search interface has two different screens: Easy and Advanced search.

Easy Search



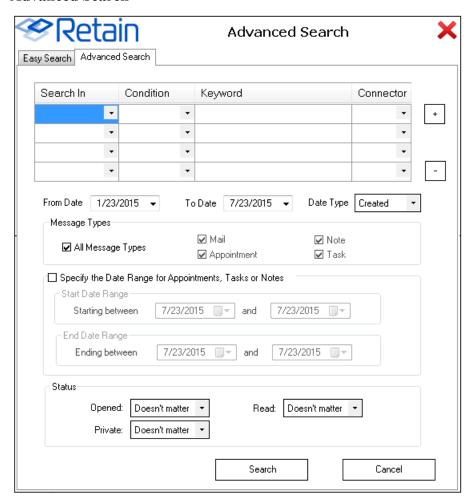
With the Easy search, the plugin only offers a simple keyword and date range. The search keywords are subject to the same restrictions and abilities as the Retain web interface. Easy search supports wildcards and quotes for search criteria. Without quotes, the search text will be searched as each word is an individual search term. With quotes, the search phrase is the criteria.

IE. Searching for ping pong will result in messages with the term 'ping', with the term 'pong', and with both, 'ping pong'

While searching for "ping pong" will only result in messages with the term 'ping pong'. Searching with wildcards: Searching for gwava.* will produce results of 'gwava.com, gwava.org, gwava.edu', and 'gwava.'.

In general, if multiple search terms are desired, move to the advanced search option. After search terms and a date range has been applied, select the 'Search' button.

Advanced Search



The advanced search allows for more control in the search process.



The 'Search in' drop-down menu allows for customization of the limiting function of the keyword.



The condition field restricts the condition of the data items in the Retain Message Store.

Finally, the Connector setting allows for the logical connection between criteria.



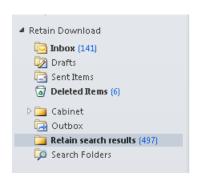
The "+ and -" buttons to the right of the window allow the user to add multiple search term lines. There is no limit to the amount of lines, or criteria, which may be added.

The Message type option allows the search to be limited to the specific selected type or types. (Note, Appointment, Task, or Mail)



The Date Range is designed with appointments, tasks, and notes. If the date range is used, the range for the desired date applies. This is essentially a date range for the date ranges. This is a good way to look for tasks and appointments which were placed into the system before, but not removed later.

The Status restriction allows the user to search for messages with a specific message state; opened, read, and private.



Once the search terms are satisfactory, select the 'Search' button to begin the search. Once the search has completed, the resulting group of messages is added to the Retain local archive, under the 'Retain search results' folder.

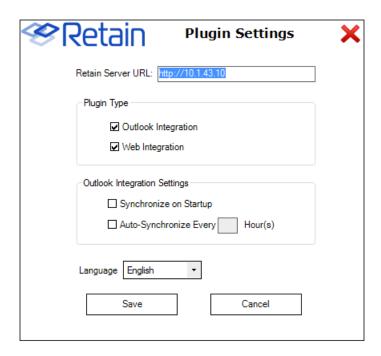
NOTE: It is important to note that the number outside the folders represents the 'unread' message count, not the total message count. To view the total message count, see the 'Items:' count at the bottom left of the Outlook window.

Quick Search

The Retain ribbon contains a quick search field, which performs a simple search of the user's archive right from the ribbon without opening any additional windows. Simply place the desired criteria in the field and click on the 'Quick Search' button. Results will be displayed the same way they are for the Easy and Advanced search. In addition, the Quick Search supports the same wildcards.

Retain Outlook Plugin Settings

The Plugin settings button opens the plugin settings window. This is where the user can modify the basic plugin settings.

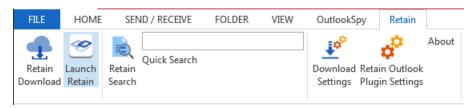


This window contains all connection and plugin settings. If the entries here are grayed-out, then the ability to change the settings has been disabled during install, and the setting information is display-only. The connection URL should be specified with just the hostname or IP address. If a port number is required, it may also be specified, but is not necessary if Apache or IIS is handling requests for the Retain Server.

Outlook as well as Web integration is controlled here, as are the download settings. The plugin may be set to automatically download messages from the archive on startup, or at a set time interval, to ensure that the plugin shows what is present in the Retain Archive.

The Language set in the Plugin Settings will change the displayed language for all Retain plugin displays.

Launch Retain



The Launch Retain button is simply a single sign-on portal to the Retain web interface for users. This plugin opens a browser and uses the account credentials and logs into the 'search messages' web interface of the Retain Server.

Silent Install Commands

The silent install commands can be utilized in deploying the plugin. The following commands are provided for administrator use during the installation.

RETAINURL : The Retain server URL

MAXSYNC : Maximum number of items to be downloaded each download session

DAYSAGOARCHIVED: Download messages older than x days ago

USERNAME : The user name of the user (email)

MESSAGESYSTEM : The type of messaging system. Values:

0 -> Exchange

1 -> GroupWise

FILTERTYPE : Filter messages based on the type during the download.

It uses bit pattern to store the values:

0 -> No filter (All types)

1 -> Mail

2 -> Appointment

4 -> Note

8 -> Task

ARCHIVEFILEPATH: The path and the file name of the desired local NST file. The path should exist but not the NST file.

SCOPE: can get two values "perMachine" and "perUser"

perMachine: The plugin will be available for all the users of the machine

perUser: The plugin will be available to a specific user.

ISWEBUI : Web UI is enabled or not. (Values: true,false)
ISNATIVE : Native plugin is enabled or not. (Values: true,false)
HASAUTOSYNC : Auto-download is enabled or not. (values: true,false)

SYNCINTERVAL : The interval of automatic downloading (integer number greater than 1, showing

the interval in hours)

SYNCSTARTUP: Automatic download enabled on startup or not (values: true,false)

CANCHANGESETTINGS: If the user can change the plugin settings (values: true,false)

CANCHANGEURL: If the user can change the Retain server URL (values: true, false)

CANCHANGEPLUGINTYPE: If the user can change the type of the plugin (values: true, false)

CANCHANGEAUTOSYNC : If the user can change the automatic download settings (values: true,false) DURATION : A non-zero positive integer showing the period which the messages are downloaded.

RANGE: The range of the period for download.

 $0 \longrightarrow Day(s)$

1 --> Week(s)

2 --> Month(s)

3 --> Year(s)

The RANGE public property along with the DURATION public property show the period of the download. For example if DURATION is set to 3 and RANGE is set to 2, they mean that the download period is for 3 month.

M+ and Netmail to Retain Migrator

System Requirements

Java JRE 8 required 4 GB RAM recommended dual-core processor or better recommended Windows 7, Server 2008 R2, Server 2012

All archives must be local to the Migrator, or located on an active mapped drive. The migrator requires direct file access.

Directories and Logs

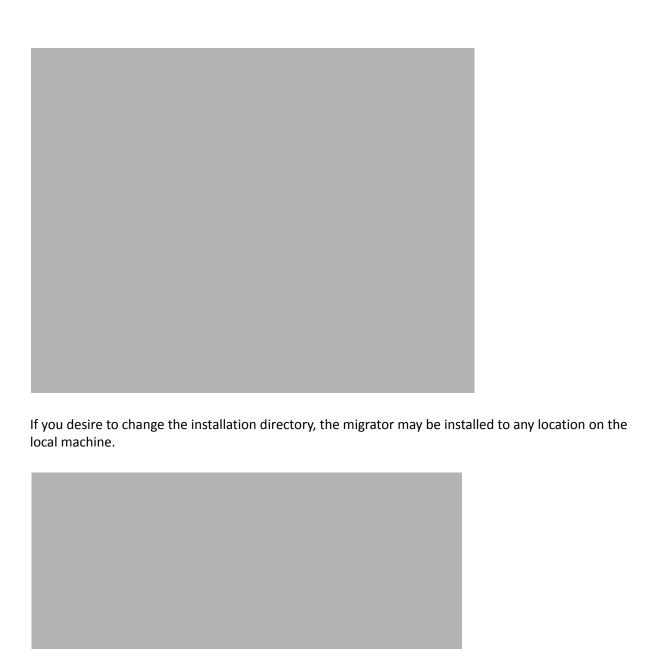
Logs are created, one for each day, in the following location by default: "C:\Users\Administrator\AppData\Roaming\GWAVA\M+ and Netmail to Retain Migrator\logs. Log locations can be changed in the Migration Context. Logs contain all errors as well as migration info for all users. After setting up the Migration Context, a .conf file can be saved anywhere and will be loaded automatically the next time the migration tool is loaded.

Install

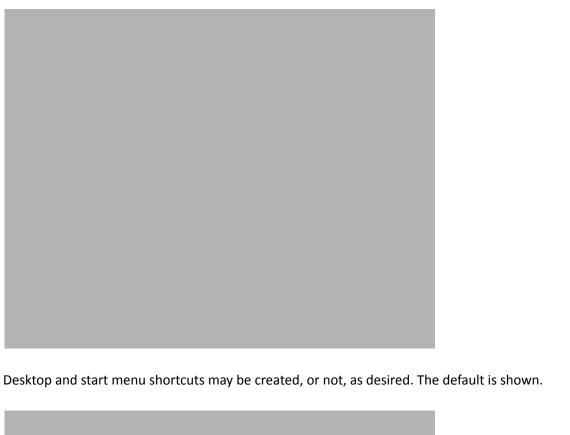
The install is straightforward. Simply execute the installation file and follow the prompts. On initial startup, if the install does not detect Java 8, then you will be notified and directed to a download for Java 8. You must restart the installation after Java 8 has been installed.



Next the system requirements are displayed.



A start menu folder will be created unless specified otherwise. The location may be modified.





All installation settings are displayed for confirmation. To confirm the settings and begin the install, click 'Next'. Use the 'Back' button to change any settings before continuing.

Once the 'Next' button has been selected, the installation will commence.



The Migrator is now ready to be run and configured on the local machine.

Quick Start Guide

Interface

The migrator is fairly straight-forward in presentation and function. In general, the migrator must be able to have access to the archives and the host mail system. The Migrator also requires an open connection to the Retain Server, and a configured Retain Worker bootstrap file which provided connection information and security access to the Retain Server. (See the Retain Administrator Guide for information on how to create a worker and download a bootstrap file.) Once configured, the migrator will connect to the host system to gather information, the local archives to gather data, and the Retain Server to migrate the data into the Retain archive. Multiple archives may be migrated at the same time. Archive size is not an issue, as the migrator only loads a single message at a time. However, due to the operations required, migrating extensive archives may take some time.

The interface has a few features, all accessed through the right-click menu. Users must be selected in the interface in order to be migrated. ONLY selected users will be migrated.

Reset - right-click on a user to reset status and errors. This will allow re-migration of problem accounts. Select All - right-click anywhere on the user panel of the migration tool to display the 'Select All' feature, to select all users listed.



Migration Process

To migrate an archive or set of archives:

- 1. Start the tool by running the executable file from the shortcut.
- 2. Configure the tool
 - Click File | Migration Context



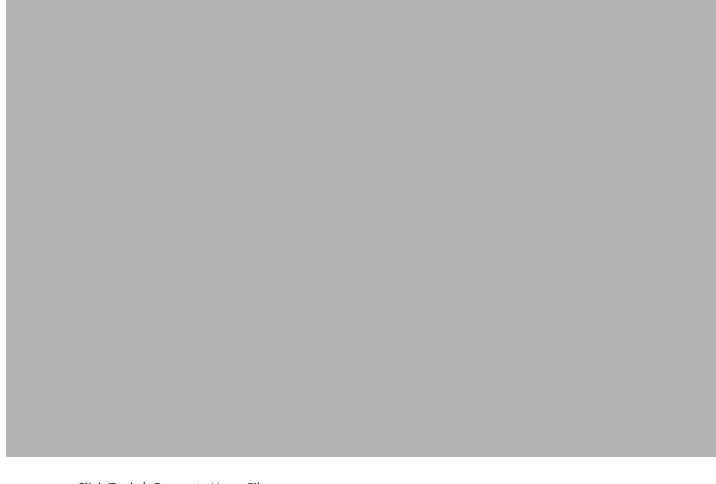


- Select the destination mail system type. (The archive type, M+ or Netmail, will be automatically detected.)
- Fill in paths to the Retain Worker Bootstrap, M+ or Netmail directory structure (Data, Audit and Attachment Folders), User List location, GroupWise SOAP connection info and # of concurrent threads
- Fill in the connection information for the mail system selected. NOTE: the connection information must be from the perspective of the migrator.
- Click Save to File and select a name and location for your configuration file
- Click Save to exit the Migration Context
 NOTE: The number of concurrent migrations may be as high as 10, or as low as 1. Performance will be determined by the processor of the local machine; a multi-core or multi-threaded processor will perform fine with multiple concurrent migrations, while machines with lower resources will be taxed by higher numbers of concurrent migrations.
- 3. Generate Users File

Users File Generation

When a user list is generated, the migration tool will attempt to locate the GUID based on the connection information in the Migration Context. If the GUID isn't found in the mail system, the archive data will be used to determine the correct id. If both methods fail, you have the option to generate a unique id by right-clicking on the user. Generating a random id should always be the last option.

WARNING: If an id is generated for a user that has already been archived in Retain under another id, duplicate accounts may be created.



- Click Tools | Generate Users File
- Select a location for the saved User List
- Click Save
- The tool will connect to the selected mail system and create a csv file of all users. The file will contain the following information:
- "userID";"userPostOffice";"userDomain";"userEmail";"userDisplayName";"user-Type";"uuid";"userFolderPath";"status"
- 4. Select Users to Migrate
 - Select users from the loaded list to migrate by left-clicking them with the mouse. Use the Shift and Ctrl keys to select multiple users
- 5. Disable Console Screen for best performance (recommended)
 - Click Tools | Show Console to toggle setting
- 6. Start the migration
 - Click the Start button at the bottom of the tool

The migration tool will display statistics on the number of successfully archived messages, number of errors, users being migrated, the status of each user (done, migrating, queued, etc.) and how long the migration for that user took.

If a Running migration is stopped, (by pressing the 'stop' button), it may be resumed simply by migrating the user again. The tool maintains a record of progress of users. If a user needs to be completely remigrated, select the user and then select 'Reset' from the right-click menu to reset the migration status. The user will then be completely re-migrated.

Retain Publisher and Viewer

Retain comes with the ability to export selected messages to a local archive for searching and viewing or to fulfill the need of a mobile archive for legal compliance. This can be extremely useful for larger systems or systems which have a high load, and where the need for review or legal compliance is being exercised, but access to the entire Retain archive is not necessary. Instead of allowing or facilitating constant access to the entire Retain Server, the Publisher can export and index groups of messages to a local archive, and the Viewer can search, view, and forward messages from the local archive.

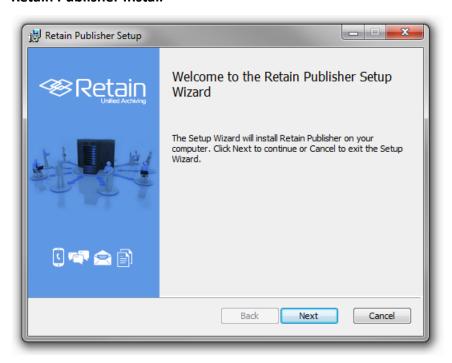
The process is performed in two parts: the Publisher and the Viewer. The publisher, using an existing account with mail export rights, connects to the Retain Server and exports the messages complying with the search request, and creates a local database archive on the host machine. The viewer accesses the local archive and allows browsing, searching, and message exportation from the local archive. Requirements:

- Network connection to Retain Server
- .Net 3.5 SP1
- Windows XP, Vista, 7, 2003, or 2008. (32-bit or 64-bit)
- PST migration requires Outlook installed. (32 and 64-bit systems are supported.)
- Retain user with rights to export mail (See User management in Retain Administration)
 NOTE: Due to file access, the publisher cannot publish archives to a network drive

The Viewer can be run or be installed as stand-alone client wherever a published archive exists, but it is recommended to install the Viewer when you install the Publisher. While the Viewer and Publisher can be installed and run separately, the Viewer must have direct access to the published archive. Do not place the published archive on a network share. Accessing the published database from a remote machine may cause instability and is insecure. For this reason, it is best to use the Viewer local to the published archive.

Download the tools from the Retain site Run the installers on the desired machine.

Retain Publisher Install

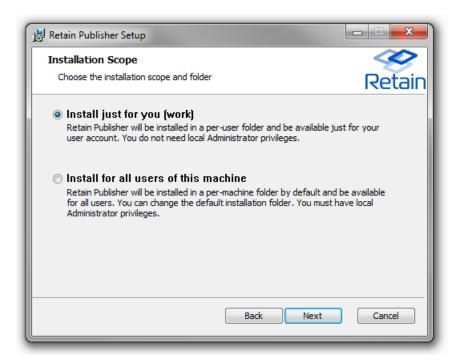


The Retain Publisher installation is very similar to the Viewer installation. Basic questions are asked, and the installer checks for .NET 3.5 SP1 before copying or installing any files. Click 'Next' to continue.

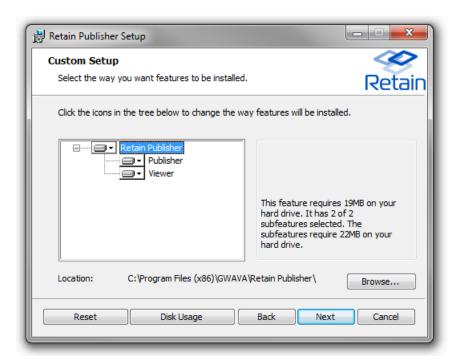


The installer checks for any programs that need to be closed for installation. Close any applications specified and click 'Next'.

Read and accept the license agreement. Select 'Next' to continue.



Select the scope of the install

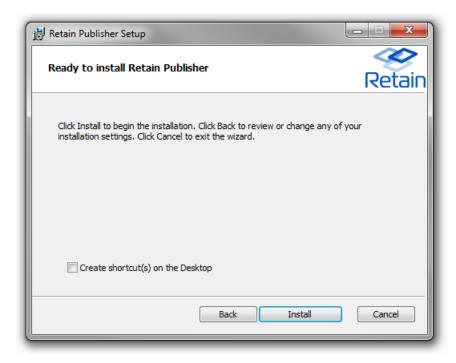


Select which components to install.

The Publisher installation provides the option to install the viewer alongside the publisher. If the viewer has already been installed, this is not necessary. It is recommended to have both the viewer on the same

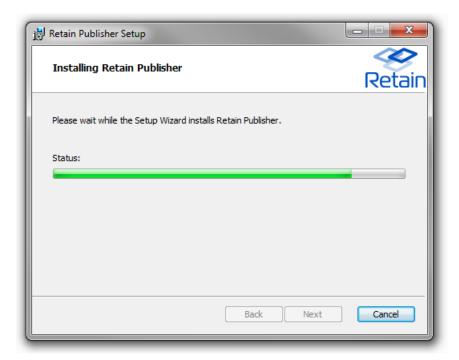
machine as the publisher. To only install the viewer, unselect the publisher. Select the desired setting and click 'Next' to continue.

Select the install location. The default is shown. If the default location does not work for the system, browse to, or specify the desired install location.

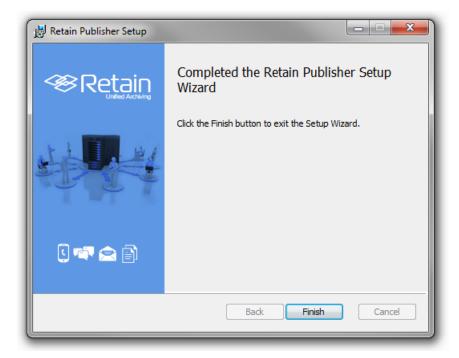


Select whether to create shortcuts or not. Default is to not create shortcuts. Select the check box to create shortcut(s) on the Desktop.

Click 'Install' to begin installation.



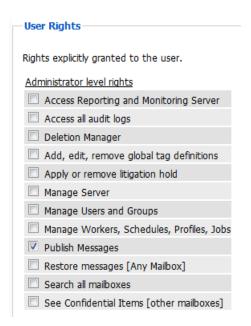
Wait for installation to complete.



After install has completed, select 'Finish'.

Retain Publisher

The Retain Publisher must be run to create the local archive for the Viewer to connect to. The Publisher does not display messages, it just accesses and exports messages into a portable message archive.



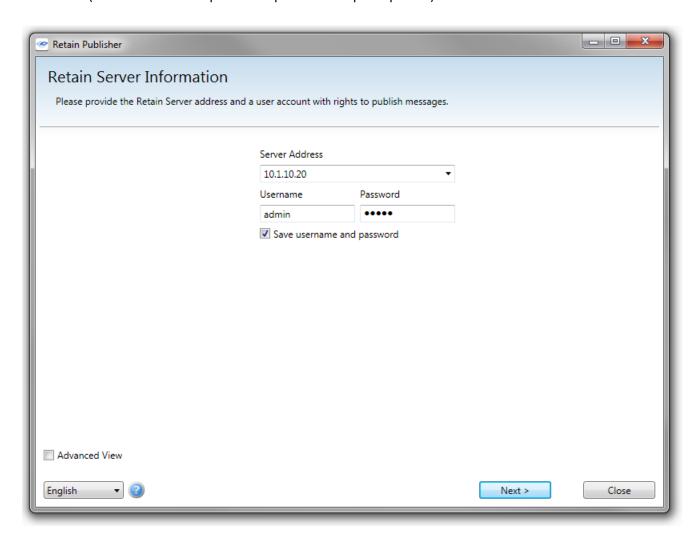
The Publisher must be used in conjunction with a user that has administrator rights to publish messages. Because the Publisher connects over the network, it requires an open network connection to the Retain Server.

The Administrator account for Retain automatically has the publish messages right, and can be used here, though it is highly recommended to create and use an auditor account with the mail export right and access to the desired mailboxes or post offices required.



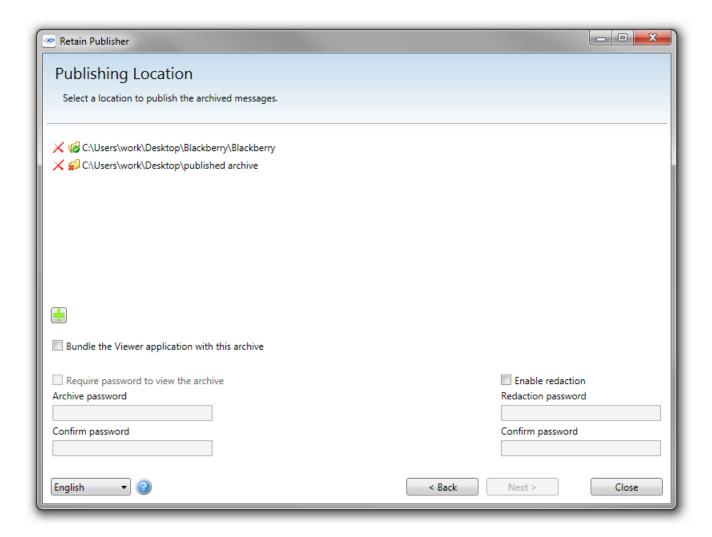
Input the Retain Server's DNS or IP Address and the login for the account with message export rights and rights to the desired mailboxes. Unless the login account has the administrator level right to 'search all

mailboxes', only the mailboxes granted to the auditor account will be accessible. Granting rights to mailboxes other than the active user's own mailbox are specified in the user rights section, under the 'mailboxes' tab. (The advanced view provides a protocol and port options.)



Select 'Next' to login to the Retain Server.

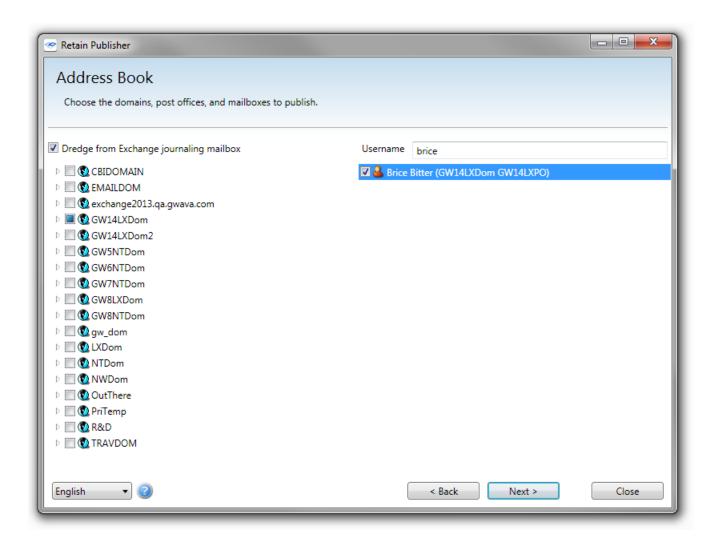
The Publisher then asks for the desired location for the exported archive messages. The messages extracted from the Retain Server will be saved in a database at this location. Select an existing location or select the green plus button and then browse to, or create a new folder for the published archive destination. If a destination is selected which already contains an existing archive, a warning of overwriting an existing archive will be displayed.



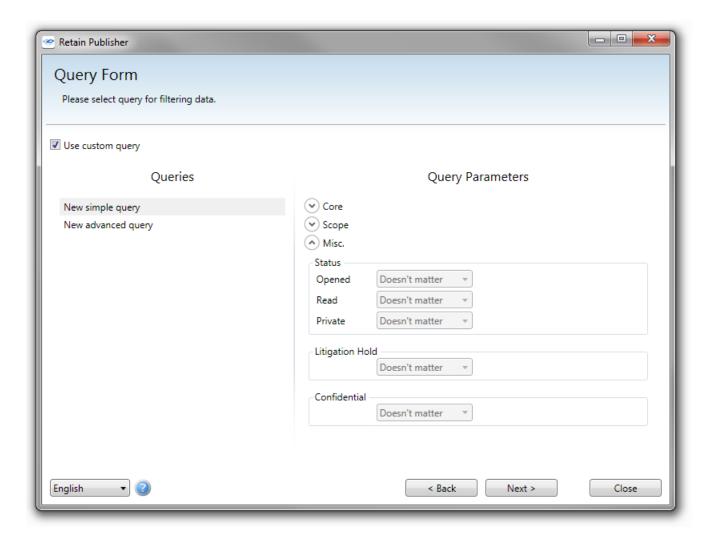
The options to password protect and use Redaction, both require passwords. To use them, select the checkbox and specify a password. The passwords can be different and both options may be present on the same archive.

When the 'Bundle the Viewer application with this archive' option is selected, the Viewer installation file is copied into the archive as well, preparing it to be completely mobile; the entire destination folder can then be copied or sent to any system with all that is required to read and search the archive. If the archive is to be distributed on a burned DVD or CD, further redaction will not function from the disk. All redaction comments should either be made before the archive is burned to a disk, or the archive should be copied from the disk to a local location where changes can be saved. When writing the archive to a read-only medium, it is always prudent to finalize the compilation.

Making the archive mobile can be very useful in situations where direct access to the Retain archive, such as during legal discovery, when email information is required to be surrendered for a user or group of users, an archive can be created and provided without disrupting current Retain operation. Click 'Next' to continue.



Select the scope of the export job. Single or multiple users, mail boxes, message servers, domains, and social media, even crossing message systems, may be added or excluded from any export job. Select 'Next' to continue.

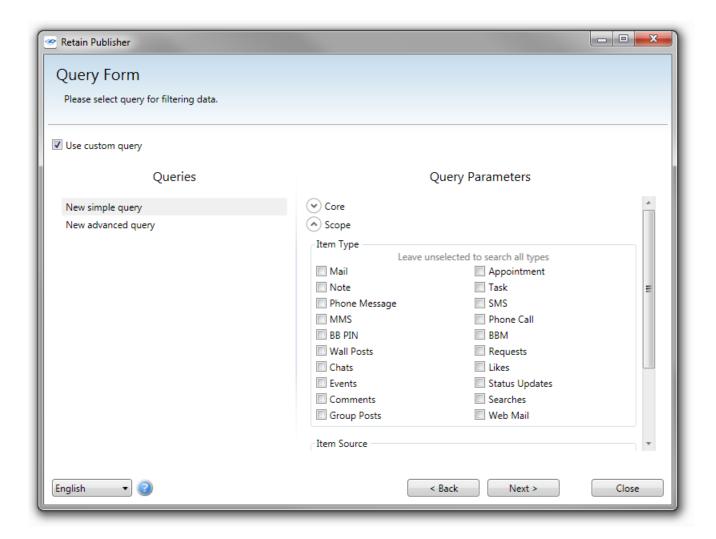


The Publisher may restrict what mail or items are published into the message archive by specifying search or query terms. If a message or item does not apply to the search terms, it will not be added to the published archive. There are three opportunities to specify search terms or criteria; Core (query), Scope, and Misc.

Each of the three different criteria areas at this step may be shown or hidden by clicking on the hide/show arrow. Select the 'Use custom query' and either the simple or advanced query options to enter search criteria. Previously saved searches or searches that have been shared with the logged-in user will be displayed in the queries menu. If the 'Use custom query' box is not checked the query parameters will remain locked and cannot be manipulated.

To add criteria to the Core query parameters, simply select the green '+' icon, then select the different desired operator, search type, and phrase.

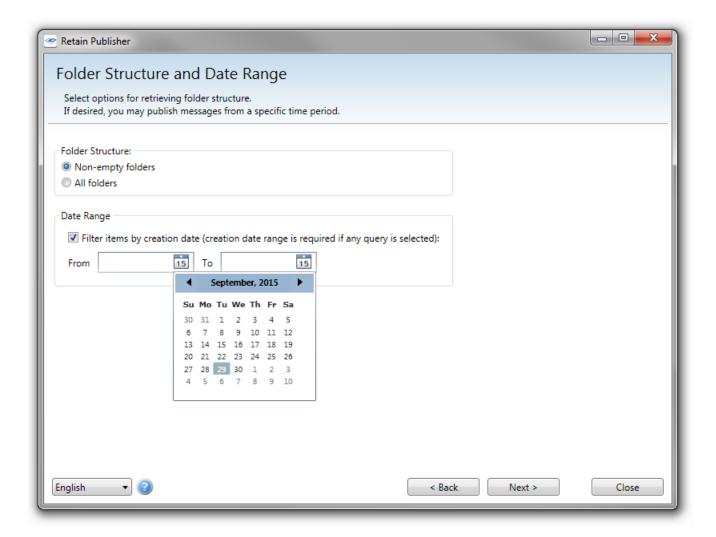
For instance, a very exclusive 'contains' search of the message contents will only publish messages that have the exact specified word or phrase in them. Fuzzy searches for approximates to the specified phrase, producing results including the words, but not in the specific order, or similar words.



The Scope details what types of messages will be searched for. If no items are checked, all are allowed, and that setting is default. To restrict to specific messages and types, (some are mail system specific), at least one item must be selected. If one item is selected, or something is desired to be excluded from the published archive, select all desired item types. The Item Type, Item Source, and Attachment size, are all independent factors and must be specified independently or left blank. A blank item source and attachment size will not restrict those criteria when combined with a specified type.

```
Any
<100 KB
>100 KB,<1 MB
>1 MB,<10 MB
>10 MB,<100 MB
>100 MB
```

Attachment sizes can be restricted or allowed into the message archive, and are acessible through the drop-down menu at the bottom of the query screen, under the 'Attachment Size' menu.



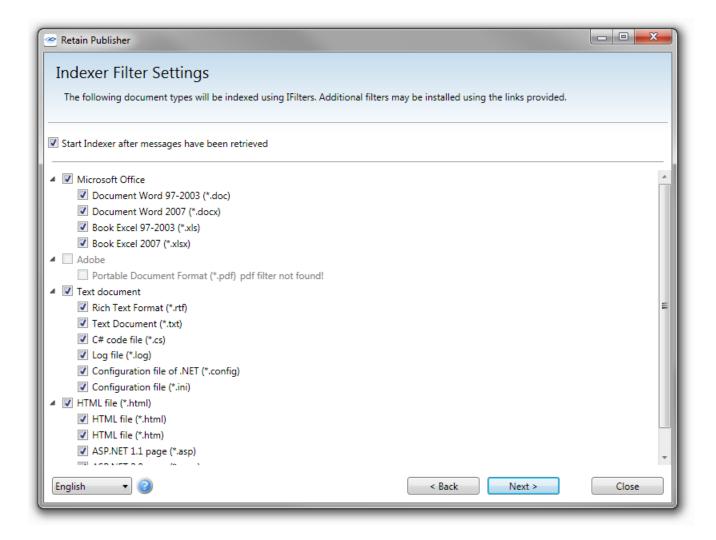
Simply selecting the desired setting will make it active in the search criteria. Set as desired and continue. The publisher can also be restricted to exclude all empty folders from the published archive.

The Date Range allows for specific item date ranges to be selected or excluded. The query will cause the publisher to only look for the specific mail required within the specified time frame. The time range may have a start date, and end date, or both.

This setting differs from the following time restriction in that the start and end times may be specified in a range, instead of a specific hard date. Only messages strictly adhering to the range will be published. If a date range is desired, enable the date range and select the desired range.

A date range is required to be specified if the active query is anything other than a previously saved query, so make sure it is set to an acceptable window for the desired data. The publisher automatically defaults the date range to one year.

Click on the 'From' or 'To' dates to reveal the data selection calendar and specify the date as desired. Click 'Next' to continue.



The Publisher can create and index the archive for faster searching and browsing by the Viewer. It is highly recommended to index a published archive. All default filters are shown. Click 'Next' to continue.

If previously published messages are present in the selected publish location then you will be presented the option to overwrite the present information, or to append the new messages to the existing message store.

The Publisher will connect to the Retain Server and export the qualifying messages from the selected user(s) mailboxes and builds the local archive.

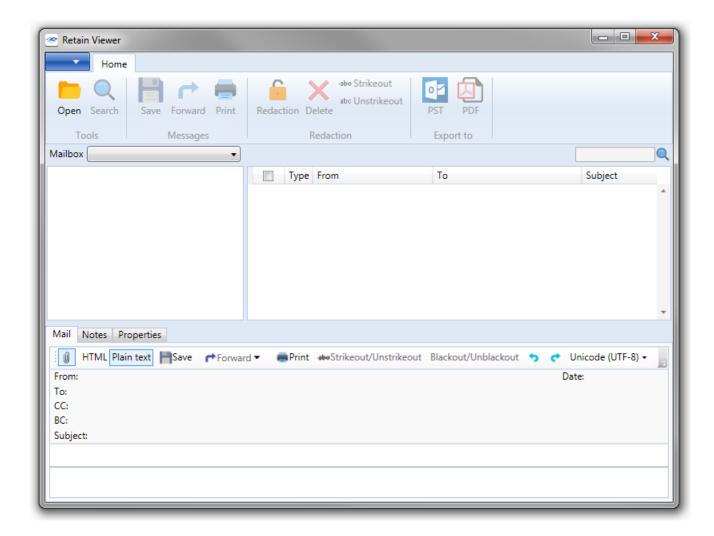
When the Publisher is done, it will display a report on the job. Depending on whether the exported archive was selected to be indexed after the job is done, the indexer will run. A .net error may occur at the end of Indexing, caused by an error in Adobe iFilter, if this error occurs, please install Adobe PDF iFilter 5.0 found on Adobe's web site.

Select 'Close' to exit the publishing wizard and the indexer will automatically run if that option was selected. With a successful publishing job, the archive is now ready to be connected to the Viewer.

If the Indexer was not automatically run, it must be run before the Viewer can search the archive.

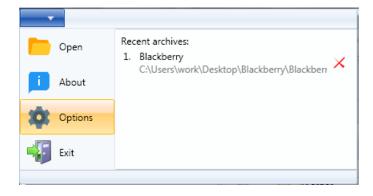
Retain Viewer

To connect to the published archive, locate the shortcut to the installed Retain Viewer and start the Viewer.

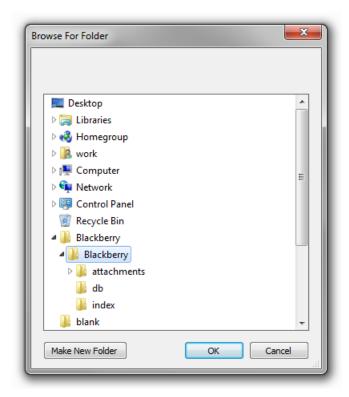


The viewer must be pointed to a valid database on initial startup. If the viewer does not prompt for the location to a published database, or if a new database is desired, select the 'Open' button and browse to and select the containing folder. If the Viewer has been previously opened, it will remember the last database opened.

The viewer is setup similar to an email client, with the mailbox and account on the left, the selected account's contents fill the space on the top right, and any selected message's contents are displayed along the bottom. Tabs allow access to the message properties or text, and highlighted buttons below the tabs determine what is shown and how, (Attachments, HTML, Plain Text). The Viewer also has the ability to forward the selected message out of the portable archive to any specified address.



To access an archive, select the 'open' button from the top toolbar, or select the 'Open' option from the 'File' dropdown menu.

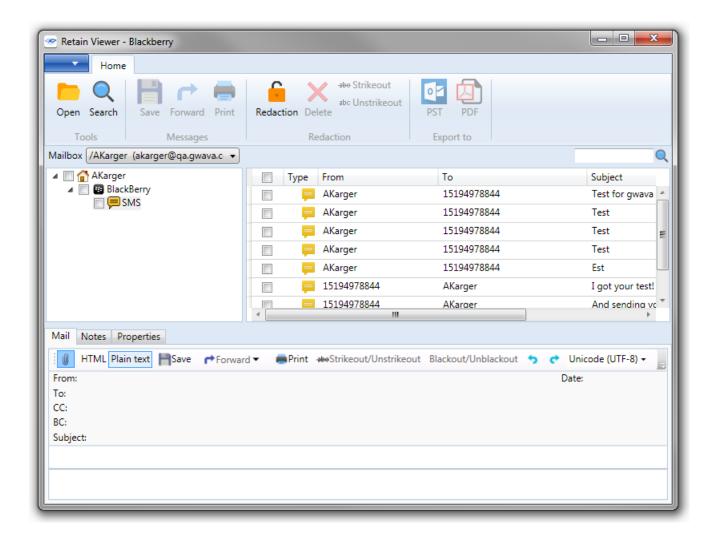


The Viewer only needs to know the base location of the archive, or the folder selected in the Publisher as the archive location. The 'Open' menu starts a browse window.

Browse to the location where the portable archive is located, select it, and click 'OK'.

If the archive was password protected, the correct password must be entered before the Viewer can gain access.

Once the Archive has opened, all mailboxes contained in the archive are displayed. Select a mailbox to access the mail in the archive.



It is important to note that while the Publisher stores time in UTC, the time displayed in the viewer, located in the 'Date' column for each message, is relative to the viewer's installed time zone, it is not in UTC. Messages published in PDF also display according local time zone. PST is exported in UTC, but Outlook rounds to the nearest minute, and displays the UTC time stamp in local time.

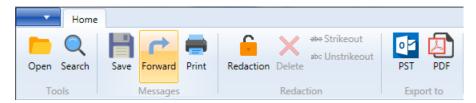


The text of the selected message will be displayed in the bottom viewer pane. The viewer contains options to show or hide the attachments, forward the message, as well as whether to view the message in plain text, or in HTML when available.



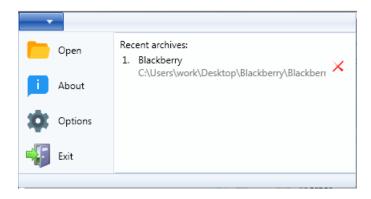
The properties tab displays the essential information on the mail item in question. The created, delivered, read status, and store date are displayed along with the identification number, message source, etc.

Forward

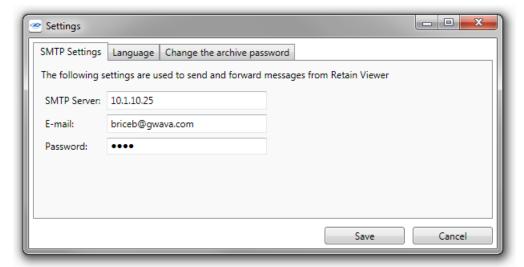


The Forward, and forward as attachment options, (found both on the main toolbar, as well as a drop down menu in the toolbar in the view window), allow the Retain Viewer to send the selected message from the archive to a destination account. 'Forward' functions exactly as the forward function in an email client. The 'Forward as attachment' creates an attachment from the selected email and attaches it to a message sent.

In order to utilize the 'Forward' option in the viewer pane, the Viewer must be configured with a mail server. To tell the Viewer which SMTP server and account to use to send messages, select 'Options' from the drop-down menu.



Settings

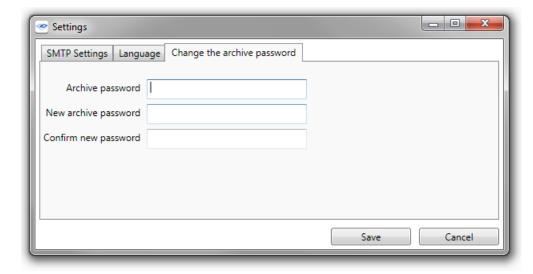


The 'Settings' configuration menu will appear. Enter the DNS or address of the SMTP Server desired to use for the Viewer system, and an appropriate account, (Username is required, password is optional), to connect and send messages.

Language and displayed document format settings may be changed as well for the specific viewer.



The 'Change the archive password' tab allows access to modify the password of the current archive. In order to change a password on an archive, the original password must be supplied first. If there is no password on an archive, a new password may be created; simply leave the current password field blank and input the desired password in the provided fields. While passwords may be changed, or created, the requirement for a password may not be removed.



Redaction



Redaction allows the viewing auditor to compile notes on the archive. The notes are appended to the entire archive, and not any individual messages, which makes Redaction extremely useful to compile messages and identify messages of interest.

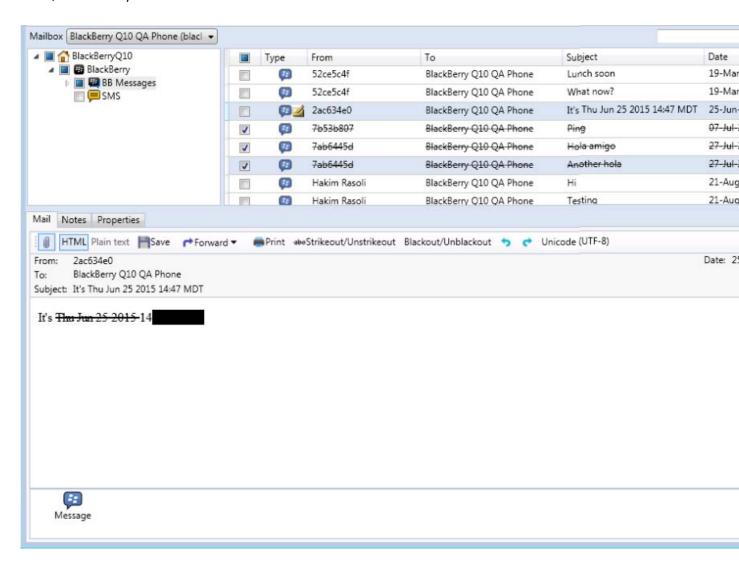


To enable redaction, the archive must first be exported with the redaction option enabled in the publisher, (if the archive was not published with Redaction enabled, the option will appear grayed out), and a redaction password must be entered.

On a redaction-enabled archive, to access the feature, it must be 'enabled' in the viewer. This option is found on the toolbar. Select 'Redaction' and enter the redaction password.

When Redaction is enabled, the 'Strikeout', 'Black out' options, and 'Notes' tab become active. The Notes tab works much like a notepad that is always connected to the archive. Items of interest or whole messages' texts can be copied and compiled in the redaction notes. Messages with notes on them are marked in the archive viewer with a notepad icon in the message list.

Redaction notes are attached to the message. Notes are only accessible for the messages which have notes, indicated by the note icon.



Strikeout/Unstrikeout and Blackout

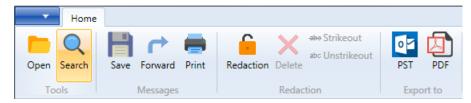
The Strikeout option enables a review of the archive before handing it over to legal scrutiny. This allows the user to denote items and messages which do not apply, or are unnecessary for the published goal. For example, if an investigation is being performed on specific activity, or internally on the account and all correspondence with a customer, the strikeout will allows messages or even sections of messages which do not apply, to be marked out.

To use the Strikeout, first Redaction must be enabled. After redaction is enabled, the options become active. To Strikeout or Unstrikeout messages in the message list, select the messages by placing a check

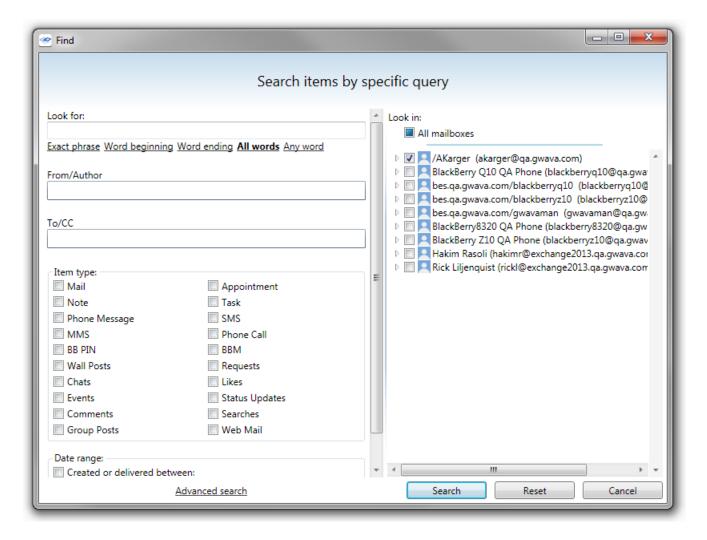
in the checkbox and then select the 'Strikeout', or 'Unstrikeout' button from the top bar. To strikeout sections in the body of the message, highlight the desired section and select the 'Strikeout/Unstrikeout' button from the view message bar.

The Blackout option allows a reviewer to remove sensitive information, such as social security numbers, from published messages. To use the Blackout option, highlight the desired text and select the 'Blackout' button. Blackout cannot be removed from text.

Searching the Archive



The Retain Viewer provides full search functionality. The search function is accessed through either the shortcut 'Ctrl+F' or by selecting 'Search' from the toolbar.

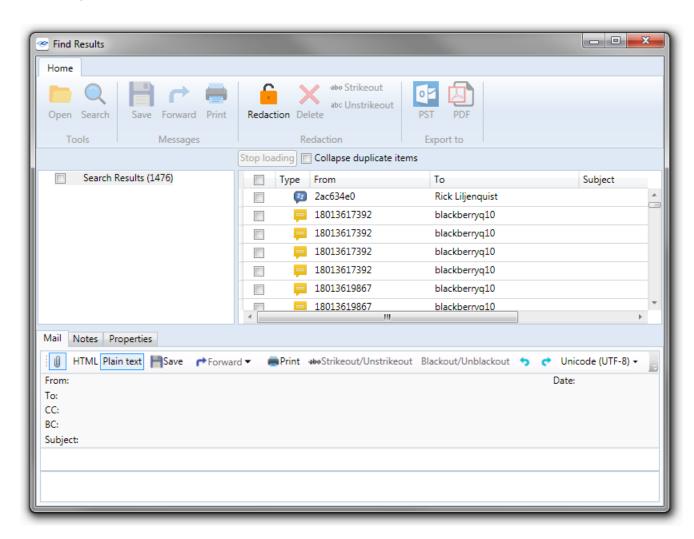


The different options for the search are shown with the criteria input below the options. Any text string or value can be searched for in the full text of the mail, or simply in the subject line. The following options are functional for both the search messages option included here, as well as the publisher search criteria. In the search 'look for' line:

- && represents 'and'
- || represents 'or'
- No operators work as an exact match search.
- For example:
- Boy girl = exact search for "Boy girl"
- Boy && girl = Boy and girl search
- Boy | | Girl = Boy OR girl search
- Boy && girl || dog && cat = Boy and girl OR dog and cat search.

NOTE: These search operators DO NOT work with the search option 'starts with', but can be utilized with all of the other base modes of the search.

NOTE: The search engine does not recognize the following English 'stop words' or articles of speech: a, an, and, are, as, at, be, but, by, for, if, in, into, is, it, no, not, of, on, or, such, that, the, their, then, there, these, they, this, to, was, will, with



These words are not indexed, and thus are not taken into account if they appear in the list of words specified for the search. This may result in some unanticipated results in "exact" searches.

In addition to the operators and text, the sending Author or the destination and carbon copy recipients of a message as well as item type and the date range can specify or restrict a search. Selection of the date range is done through an interactive calendar.

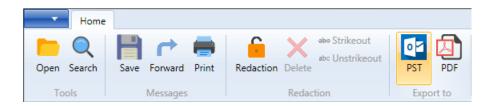
Be sure to select the desired mailbox or mailboxes to apply the search through. (The depicted example archive only contains one mailbox.) Select 'Search' to begin.

When the search is finished, the results will be displayed in a separate results window, identical to the main viewer interface.

PDF and PST Export

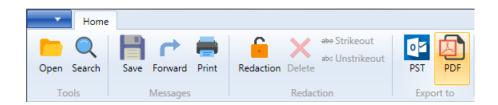
Retain Viewer can export selected files and messages from a mailbox or archive to a PDF or PST archive file. NOTE: PST export requires Outlook, or the associated plugin. The viewer will prompt for plugin installation if necessary. It is highly recommended to have Outlook installed.

To migrate to PST, select desired files in the view window then select 'PST' from the toolbar.



A window confirming the location and file will open. When the 'Save' button is selected the PST file is created and available for use.

The Retain Viewer can also save selected messages as PDF. Select the desired messages and click 'PDF' in the toolbar.

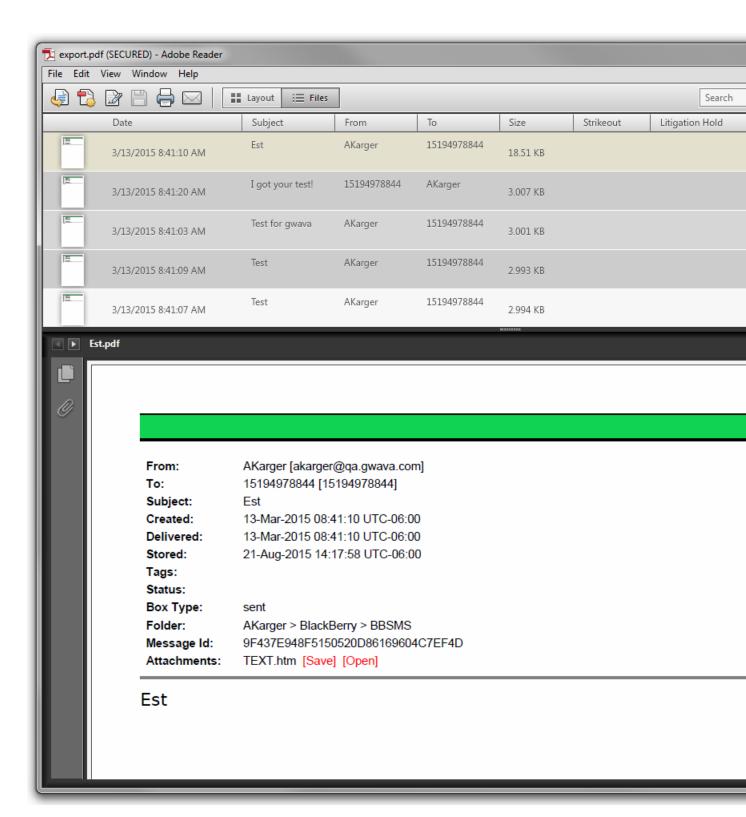




The selected messages are then published in a single PDF. In addition to the standard file name and location, the PDF has title, exported by, and comment sections that the user is prompted to specify before the viewer publishes the PDF.

The finished PDF contains all the information displayed in the viewer, as well as an option to save individual text or message sources. (This information includes redactions, litigation tags, confidential tags, strikeouts etc. from the Retain Server archive as well as the Viewer.)

The published PDF displays the separate emails along the top in a list and displays the message information in a lower pane, as shown below.



Retain PST Importer

The Retain PST Importer is a tool which inputs mail from PST archives into the Retain system. The PST Importer is a utility which is installed on a Windows machine, local to the PST archive and works in coordination with the Retain Server.

System Requirements

Java Runtime Environment 8 (No other Java can be present on the system.) Windows 7, Windows 2008 R2

Open and direct network connection to the Retain Server

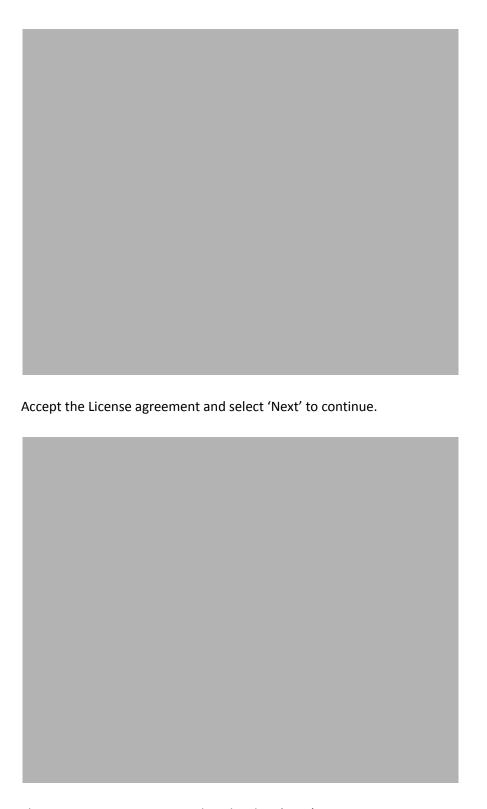
NOTE: If the correct Java is not present in the system, the Importer install will not start, but will provide a link to the download of the appropriate JDK. Select the correct version, (32 or 64-bit), download and install it. Due to the way Windows interacts with Java, there can be no other version of Java on the host machine. Uninstall any other version.

Installation

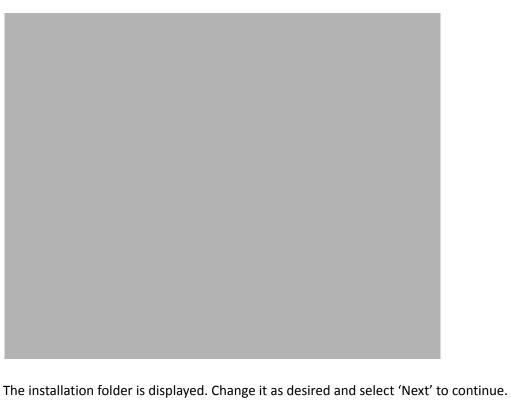
Installation of the PST Importer is simple. Locate the 'Retain PST Importer Install.exe' file and run it.



Once initialized, the installation will display this page. Select 'Next' to continue.

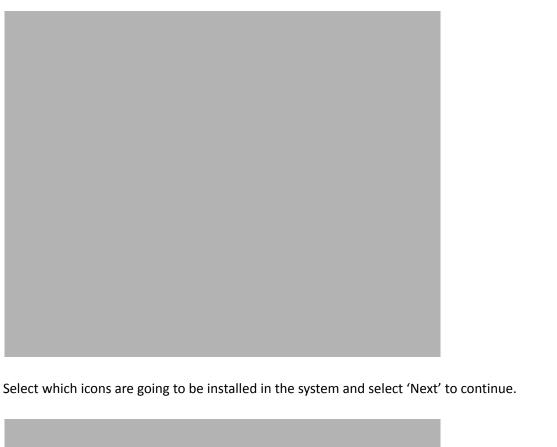


The System Requirements are listed. Select 'Next' to continue.





The PST Importer may be installed with a shortcut in any desired start menu folder. The default folder is displayed. Change as desired and select 'Next' to continue.





Review the installation settings. If any settings need to be changed or revised, use the 'Back' button to return to the appropriate pages and change them. When correct, select 'Install' to install the program.



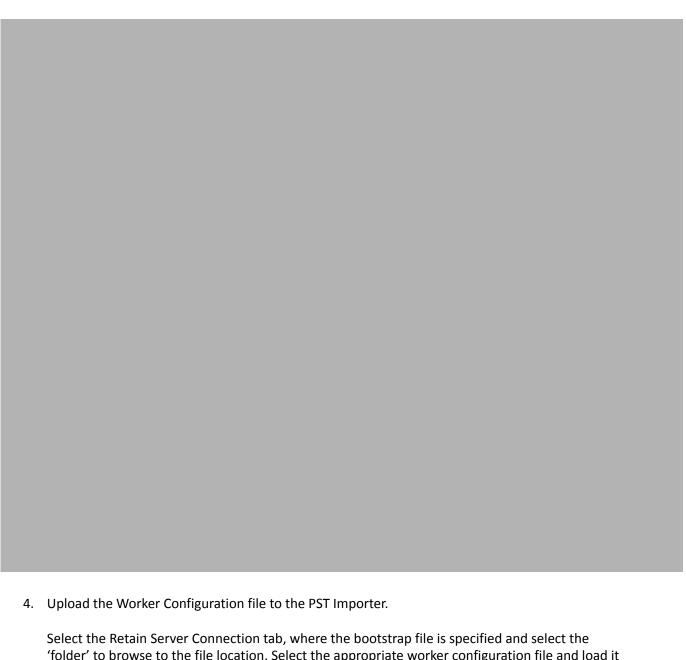
Once the installation is completed successfully, select the 'Finish' button.

Configuration and Use

The PST Importer requires a Worker configuration and a bootstrap file to know where and how to communicate with the Retain Server. A new worker must be created, configured, and the configuration file loaded into the PST Importer. Follow the steps below.

1. In the Retain Server administration console, create a new Worker. It would be appropriate to name it 'Importer' or 'pst importer' for ease of remembrance, but any name desired will do. Remember that this worker should never be configured with any jobs in the Retain Server. A general configuration is all that is required; no specific settings need to be changed from default, except the connection information if it needs to be modified in accordance to the needs of the local network. It is imperative that the PST Importer has an open connection to the Retain server. When the settings are correct, save the new Worker.

2.	Download the bootstrap file.
3.	Run the PST Importer utility. Locate the Retain PST Importer icon or start menu item and run the
J.	program.
	On initial startup, the Retain PST Importer opens to the configuration page. The worker configuration file contains all the connection information for the Retain Server. (If the location changes,
	the worker configuration file will need to be re-downloaded.)



'folder' to browse to the file location. Select the appropriate worker configuration file and load it into the Import tool.

Use

Use of the PST Importer is very simple. The PST Importer tool may be restricted to import only messages in the archive which fall into a specified date range or it will attempt to import the entire archive. The PST Importer may be restricted to import only messages which meet selected criteria:

- A specified date range (inclusive)
- Before a specified date (exclusive)
- After a specified date (exclusive)



If desired, the PST Importer may send a notification message when the job is complete. The email will contain a job summary including:

- The total number of messages selected for import
- number of successfully process messages
- number of messages not processed



To receive a notification email, an SMTP connector needs to be specified and settings provided. Security protocol options are: 'Unencrypted, TLS, and SSL'. Test the connection to ensure the settings are correct.

When the desired archive or archive have been added and settings have been configured, select the 'Start' button to begin importing the archive. Rolling information is displayed in the information window at the bottom of the screen.

To add or remove an archive to or from the PST Importer tool, use the 'add' and 'remove all' buttons at the top of the archive window.

Adding an archive opens a browse window to allow selection of the archive. Adding an archive requires the selection of the desired messaging system, (select the system the user will be using so messages will be added to their archive). The user can select the file name pattern corresponding to the PST archives and whether to import the archive into existing root folders or to have the tool import the archive in to automatically created subfolders.

Users can import messages from the PST archive for either existing users in Retain, or if the user does not currently exist in Retain, it will be automatically created.



When complete, select 'Browse' button to browse to and select the archive. To remove all the archives from the list, select 'remove all'.

Logs

Log files are generated and placed in the following default location: C:/Users/<username>/AppData/Roaming/GWAVA/Retain PST Importer

GroupWise WebAccess Plugin

Overview

The Retain for GroupWise WebAccess Plugin is a quick and effective way to integrate access to the Retain archive into WebAccess for users. The Plugin adds a button which will automatically take users to their personal archive. The archive interface will be displayed inside the WebAccess window, and provides all the same access and management of the archive which users can experience by directly logging in to Retain.

Supported Software

GroupWise WebAccess 8x, 2012, 2014

System Requirements

Java 1.7 or later

Functioning GroupWise WebAccess

Functioning Retain Server, with configured GroupWise module

Open network connection to the Retain Server

If SSL is to be used on the web server, either IIS or Apache, then SSL must be enabled on the Retain server as well.

Installation

To install the plugin, execute the installation file.

Linux:

Linux requires root access to install the plugin. Change the files to executable and run the files.

- chmod +x *.sh
- ./installgui.sh (if a graphical install is desired)

or

• ./installtext.sh (if a text install is desired)

Windows:

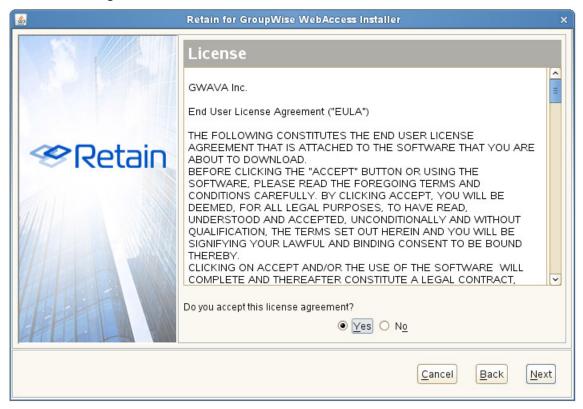
On Windows, double click and run the .bat file.

install.bat

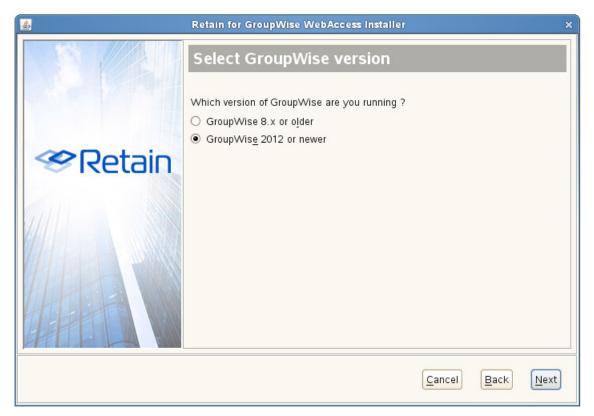
Graphical install shown.



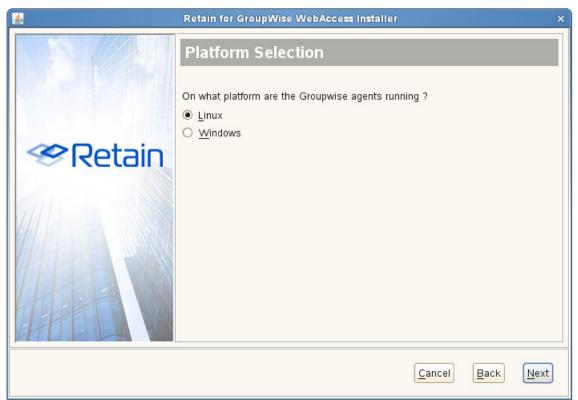
Click 'Next' to begin the installation



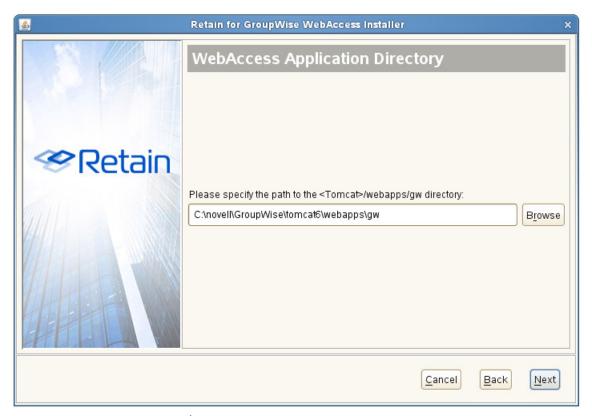
Accept the license agreement.



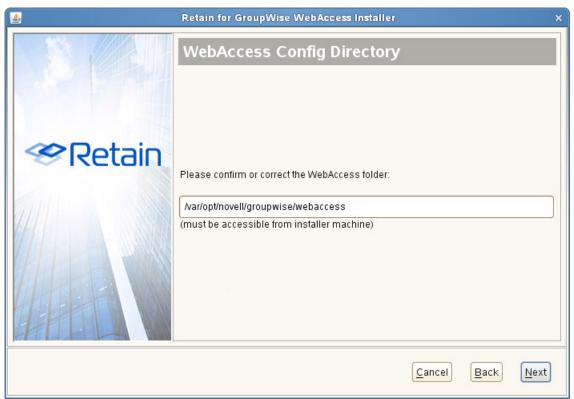
Select the GroupWise WebAccess version.



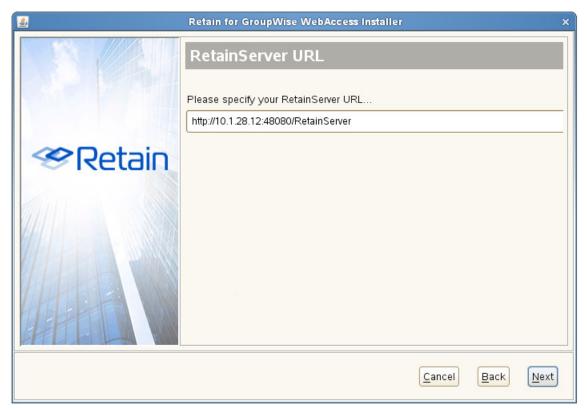
Select the platform the WebAccess is being installed to.



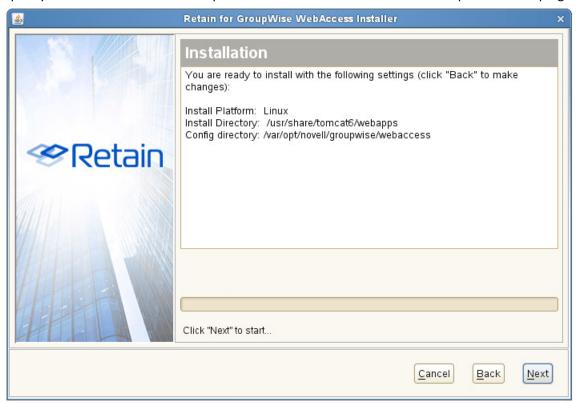
Specify the active ...webapps/gw directory, or browse to it.



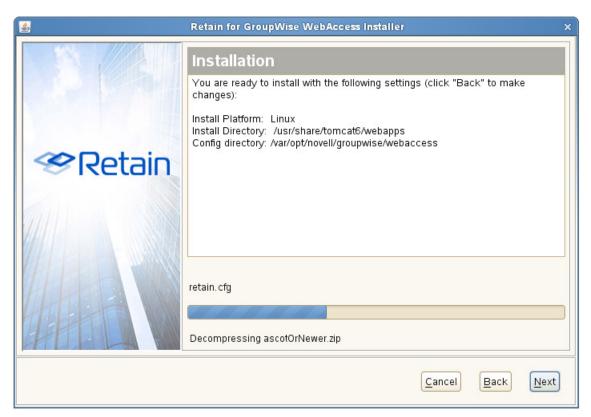
Specify or confirm the WebAccess configuration directory.



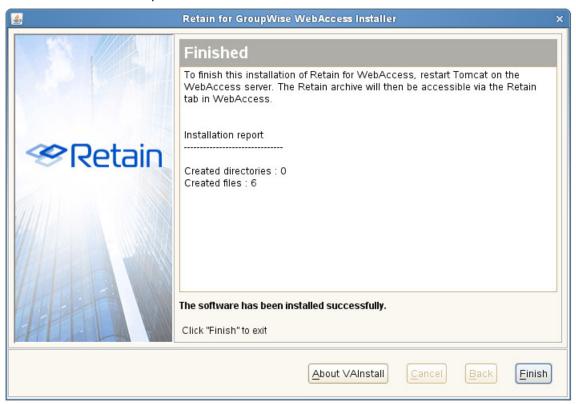
Specify the Retain Server URL. An open connection to the Retain Server is required for the plugin to work.



Confirm the settings. Click 'Next' to begin the install, or 'Back' to change the settings.



Wait until the install completes.

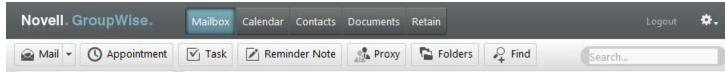


Once finished, click 'Finish' to exit the installer. The installation will not fully complete until Tomcat is restarted and the plugin files are loaded by the web server.

The plugin is now installed. Once WebAccess has been restarted, simply login as usual.



After logging in, users will see the interface they usually see, with an additional 'Retain' button on the end of the top ribbon.



For users to access their own archive, all they need to do is click on the 'Retain' button and their archive will automatically open inside the WebAccess window.

GroupWise Archive Migration Tool

GroupWise Archive Migration Tool

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

There are two modes:

- 1. Run at all workstations mode.
 - From a command line, you'll run this and it will log into the users mailbox, and then open the local archive. It will then feed messages directly to Retain.
- 2. Centralized run.

This only works if you have all archives stored at a central point.

In this case, one can conveniently run all of the migrations from a single workstation, and not involve the users at all.

System Requirements

- A Windows workstation running GroupWise client 7.0.1 or 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 Group-Wise client and sends it to a Retain Server.

The migration tool reads the path to the archives as defined in the GroupWise mailbox that you're accessing. That path must be accessible by the workstation you're on. It will then ask GroupWise to open the archives and it will read the items in the archive and send them to the Retain Server you have 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

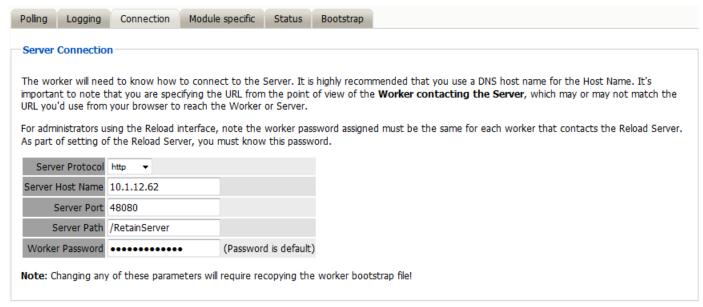
The GroupWise client on the workstation you're using must be able to open the archives you wish to migrate. The process depends on GroupWise being able to access these archives. If you're migrating the archives of another user, as you would be doing in multi-user mode, you must be able to access that 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.

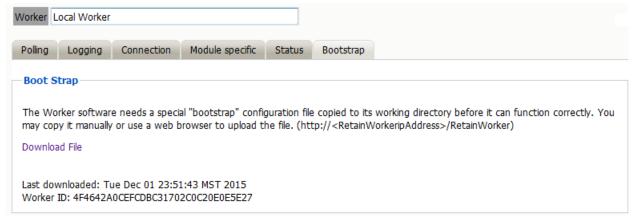
Setting It All Up

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.

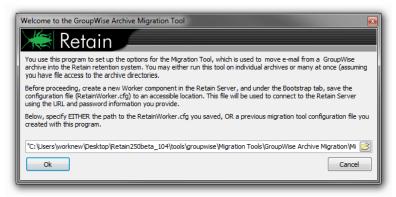
- 1. Copy the migration tool files to any directory accessible to your workstation
- 2. Log in to the Retain Server you want to migrate the archives to Create a new worker. You're not actually installing a new Retain Worker anywhere; you're creating a new bootstrap file for the use of the migration tool 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



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



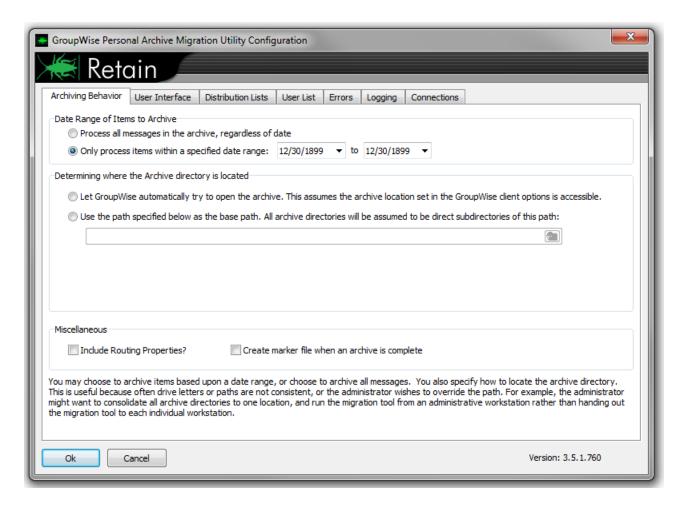
4. Now start up the migration tool configuration utility. A window will open on top of the configuration utility requiring the location to find the bootstrap file you have just saved.



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

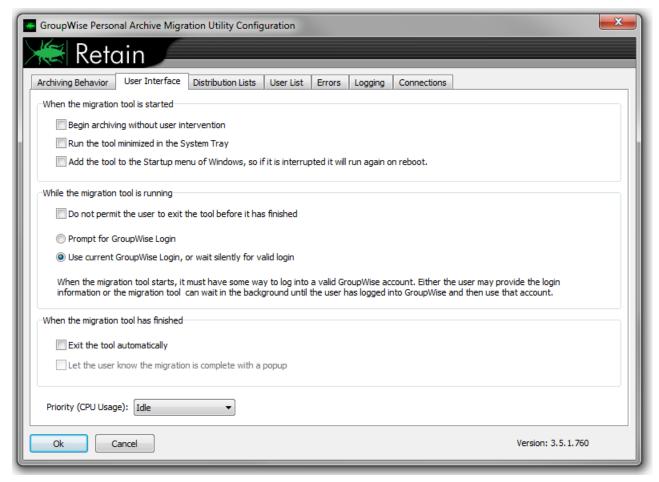
Archiving Behavior

- Date Range limits or specifies specific mail to process based on a date range.
- Archive Directory The location of the archive or archives to be migrated. The default path is assumed unless selected to be specified.
- Miscellaneous contains the options to preserve the routing properties in the archive or to add a marker file when the migration is complete.



User Interface

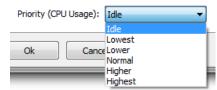
When the Migration Tool is Started - Begin without user intervention the archive process immediately begins without any prompts. Run minimized sets the migration utility to run the migration in the background while the workstation completes other work. The tool may also be set to run automatically on system boot, in case something interrupts the archive job.



When Migration Tool is Running - Contains options to prohibit the user exiting the program before the archive migration is complete and GroupWise login settings. The tool must be able to login and gain access to the archives.

When the Migration Tool has Finished – Contains the behavior of the migration utility on completion; automatic exit, notification, or neither.

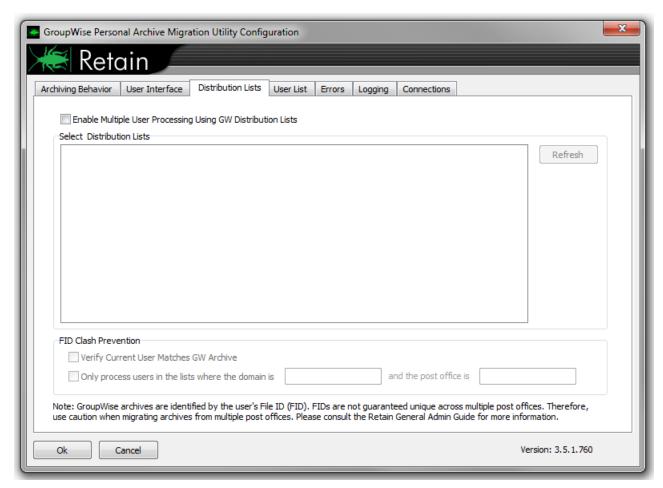
CPU priority usage is also set here.



Distribution Lists

In order to process more than one user archive the option must be enabled at the top of this page and the users must be verified against the GroupWise system.

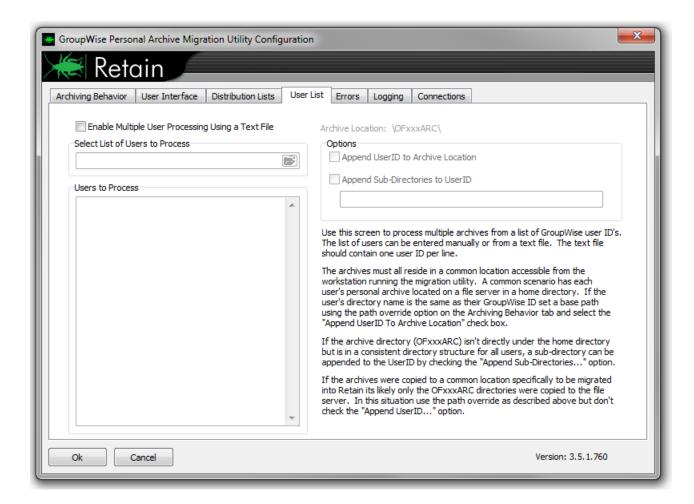
It is recommended to only process multiple archives from the same post office.



The Distribution Lists will only be populated if it is enabled and if the 'Refresh' button is selected. It displays the available distribution lists. To have the connection work correctly, the GroupWise Client Parameters on the 'Connections' tab must be filled-out with the GroupWise Server IP and port. Displayed lists can be selected. The archives belonging to the users in the selected distribution list(s) will be migrated. If users from multiple Post Offices are to be migrated, employ some kind of FID clash protection by verifying archives or only processing users which can be verified by a specified Domain and Post Office, (specified in DNS:port or IP address:port).

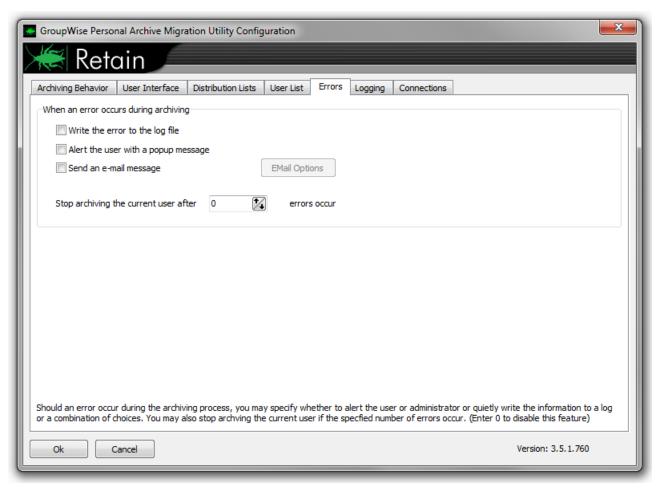
User List

The User List screen allows the migration of multiple users specified by a user list. The user list is a plain text file with one user ID per line. To enable the process, select the 'Enable Multiple User Processing Using a Text File checkbox, then browse to and select the desired text file.

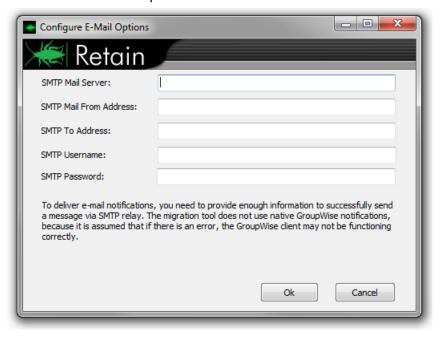


Errors

The error limit and actions taken when an error is encountered are specified on this tab.



Errors may be written to the log file, (specified in the next tab), sent in a pop-up message to the user, or sent as an email to the specified address.



The limit of how many errors may be encountered before the migration utility quits the current user is specified. A setting of '0' sets the limit to infinite.

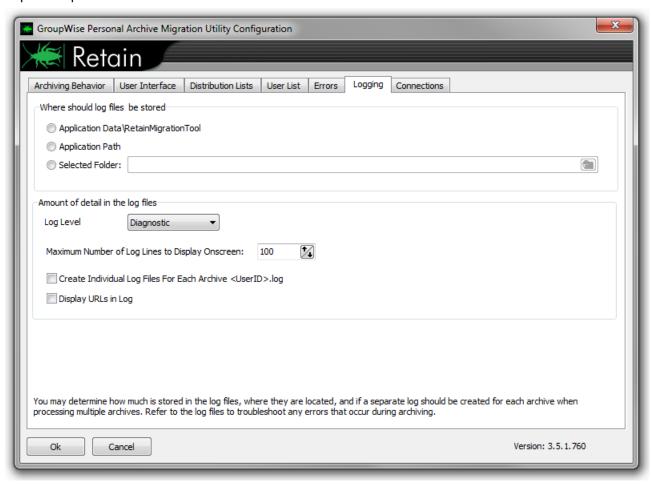
Important: By default, the current user will be skipped after 20 errors by default.

If an email is desired for each error encountered, email settings must be configured. Select 'Email Options 'and input the connection information, source address, destination address, and user login information for the SMTP Mail Server.

Logging

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

Where Should Log Files be Stored - A specified path, the path to the current location of the application, or a specified path must be selected.

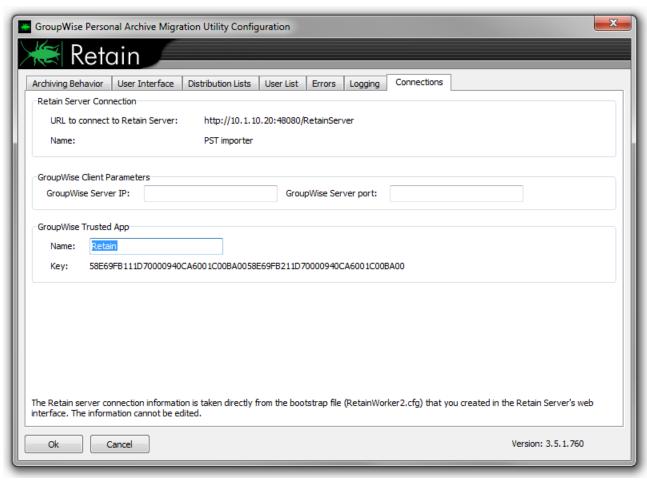


Amount of Detail in the Log Files – The logging level determines how much information is provided in the log files. Unless troubleshooting, log levels of 'normal' and 'errors only' are sufficient. The buffer size for the logging screen determines how much history the running log screen of the migration tool contains.

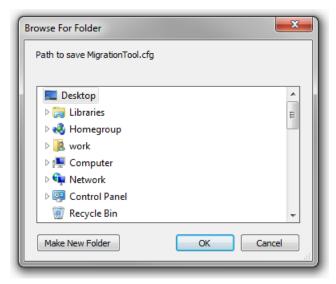
Connections

This is an informational page showing the connection settings from the RetainWorker.cfg created in the Retain Server. If this information is incorrect, the Retain Server Connection settings must be modified under the worker configuration in the Retain Server, and a new bootstrap file downloaded for use.

The GroupWise Client Parameters contains the connection information used by the Distribution Lists configuration to allow the Migrator to connect to the GroupWise system.



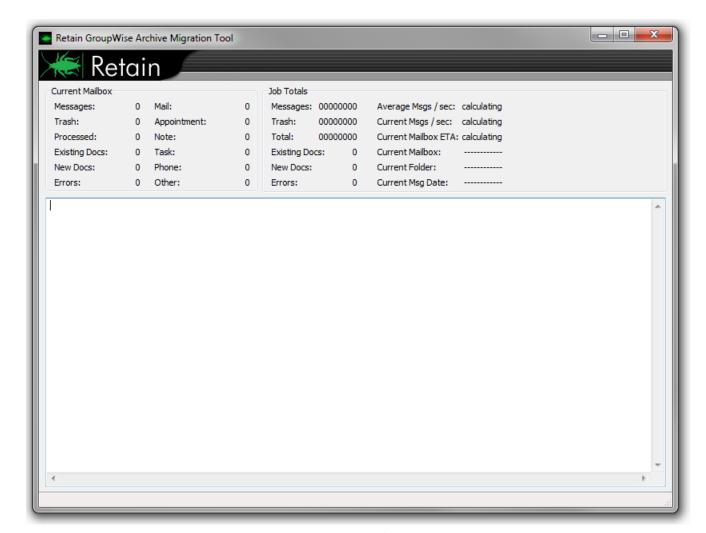
Once the configuration has been setup as desired, click 'Ok' and select the location to save the changes into a 'MigrationTool.cfg' file.



Running the Migration

The migration tool will use the configuration file to run the job as you have specified in the configuration tool. You may run the configuration tool as often as you need to set up the migration job to your liking.

When run, the tool will prompt for the configuration file, or automatically load the configuration file if the configuration is found it the same directory, and immediately begin the migration. When the tool runs, you will see a screen like this:



You will see any errors in this screen and it will display a notification when the job is complete.

For each user in a multi-user job, it will read the location of the archives from GroupWise and then it will attempt to open the archives. It is critically important that the workstation used for the migration can access the archive files.

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.

Retain Component Licensing

Copyright 2015 © GWAVA Inc. All rights reserved.

GWAVA Retain uses many components covered under the following licenses. The licenses follow the listing below.

License Components

Apache 2.0 License Apache, Open Symphony Quartz, Apache Log4j & Log4j Additions, Apache Derby, Apache XML-RP

53 Toolkit, SPring Framework

LGPL EDS, Hibernate, C3PO Connection Pooling, Jericho HTML Parser, Textmining, 7-zip

BSD License XStream, JDOM, MetaStuff DOM4J, Object Web ASM, ANTLR

BOuncy Castle APIS Special license included below PDFBox Special license included below

Open Software openXML Indexer

License v.3.0

Mozilla Public iText

License

GNU GPL V. 2 VAInstall

Bitstream Vera Fonts DejaVuSans Fonts

Copyright

ICU Licence Java Utilities

Apache 2.0 License

Apache License

Version 2.0, January 2004

http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

- Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
- Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
- Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - Tou must give any other recipients of the Work or Derivative Works a copy of this License; and
 - Ω $\mathfrak D$ You must cause any modified files to carry prominent notices stating that You changed the files; and
 - You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, exclud-

ing those notices that do not pertain to any part of the Derivative Works; and ♀ ① If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

- Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
- Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
- Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITH-OUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
- Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
- Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by rea-

son of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at

http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

LGPL - GNU Lesser General Public License

GNU LESSER GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. http://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

As used herein, "this License" refers to version 3 of the GNU Lesser General Public License, and the "GNU GPL" refers to version 3 of the GNU General Public License.

"The Library" refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An "Application" is any work that makes use of an interface provided by the Library, but which is not otherwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of

using an interface provided by the Library.

A "Combined Work" is a work produced by combining or linking an Application with the Library. The particular version of the Library with which the Combined Work was made is also called the "Linked Version".

The "Minimal Corresponding Source" for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isolation, are based on the Application, and not on the Linked Version.

The "Corresponding Application Code" for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

☐ ☐ <a>© Exception to Section 3 of the GNU GPL.

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

Conveying Modified Versions.

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

- ① Under this License, provided that you make a good faith effort to ensure that, in the event an application does not supply the function or data, the facility still operates, and performs whatever part of its purpose remains meaningful, or
- \mathfrak{Q} \mathbb{Q} under the GNU GPL, with none of the additional permissions of this License applicable to that copy.
- ☐ I Object Code Incorporating Material from Library Header Files.

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

- © Give prominent notice with each copy of the object code that the Library is used in it and that the Library and its use are covered by this License.
- A DAccompany the object code with a copy of the GNU GPL and this license document.
- Combined Works.

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:

- © Tigory Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License.
- \mathfrak{Q} \mathfrak{D} Accompany the Combined Work with a copy of the GNU GPL and this license document.
- To For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document.
- Ω Do one of the following:
 - ©Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the

user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.

- Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (a) uses at run time a copy of the Library already present on the user's computer system, and (b) will operate properly with a modified version of the Library that is interface-compatible with the Linked Version.
- M. Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such information is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option 4d0, the Installation Information must accompany the Minimal Corresponding Source and Corresponding Application Code. If you use option 4d1, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)
- Combined Libraries.

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

- ② Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License.
- \mathfrak{O} ①Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.
- Revised Versions of the GNU Lesser General Public License.

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a certain numbered version of the GNU Lesser General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that published version or of any later version published by the Free Software Foundation. If the Library as you received it does not specify a version number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall apply, that proxy's public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.

BSD License – Individual Component Licenses

ANTLR 3 License

[The BSD License] Copyright (c) 2003-2007, Terence Parr All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the author nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

ASM

Copyright (c) Object Web, ASM project.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the ASM Project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

DOM4J

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain	in copyright statements and notices.	Redistributions
must also contain a copy of this document.		

Redistributions in binary form must reproduce the above copyright notice, this list of conditions

- and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name "DOM4J" must not be used to endorse or promote products derived from this Software without prior written permission of MetaStuff, Ltd. For written permission, please contact dom4j-info@metastuff.com.
- Products derived from this Software may not be called "DOM4J" nor may "DOM4J" appear in their names without prior written permission of MetaStuff, Ltd. DOM4J is a registered trademark of MetaStuff, Ltd.

Due credit should be given to the DOM4J Project - http://www.dom4j.org

THIS SOFTWARE IS PROVIDED BY METASTUFF, LTD. AND CONTRIBUTORS "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL METASTUFF, LTD. OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright 2001-2005 (C) MetaStuff, Ltd. All Rights Reserved.

JDOM

Copyright (C) 2000-2004 Jason Hunter & Brett McLaughlin. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions, and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions, and the disclaimer that follows these conditions in the documentation and/or other materials provided with the distribution.
- The name "JDOM" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact <

request@idom.org>.

Products derived from this software may not be called "JDOM", nor may "JDOM" appear in their name, without prior written permission from the JDOM Project Management < request@jdom.org.

In addition, we request (but do not require) that you include in the end-user documentation provided with the redistribution and/or in the software itself an acknowledgement equivalent to the following:

"This product includes software developed by the JDOM Project (http://www.jdom.org/)."

Alternatively, the acknowledgment may be graphical using the logos available at http://

www.jdom.org/images/logos.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE JDOM AUTHORS OR THE PROJECT CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the JDOM Project and was originally created by Jason Hunter < <u>ihunter@jdom.org</u> > and Brett McLaughlin < <u>brett@jdom.org</u> >. For more information on the JDOM Project, please see < <u>http://www.jdom.org/</u> >.

XStream

Copyright (c) 2003-2006, Joe Walnes Copyright (c) 2006-2007, XStream Committers All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of XStream nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

PAL Artistic License

The Artistic License

Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less cus-

tomary fashion, plus the right to make reasonable modifications.

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

- You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
- You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
- You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
 - Dplace your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as ftp.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
 - \mathfrak{Q} ① use the modified Package only within your corporation or organization.
 - Trename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
 - <u>□</u> ① make other distribution arrangements with the Copyright Holder.
- Tou may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
 - ① distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
 - Ω ①accompany the distribution with the machine-readable source of the Package with your modifications.
 - ①accompany any non-standard executables with their corresponding Standard Version executables, giving the non-standard executables non-standard names, and clearly documenting the differences in manual pages (or equivalent), together with

instructions on where to get the Standard Version.

- <u>□</u> ① make other distribution arrangements with the Copyright Holder.
- You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own.
- The scripts and library files supplied as input to or produced as output from the programs of this package do not automatically fall under the copyright of this Package, but belong to whomever generated them, and may be sold commercially, and may be aggregated with this Package.
- © C or perl subroutines supplied by you and linked into this Package shall not be considered part of this Package.
 - The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.
 - THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The End

Bouncy Castle APIs

Bouncy Castle

Copyright (c) 2000 - 2006 The Legion Of The Bouncy Castle (http://www.bouncycastle.org)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

PDF Box

PDFBOX

Copyright (c) 2003-2005, www.pdfbox.org All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of pdfbox; nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

XFire

Copyright (c) 2005 Envoi Solutions LLC

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Jaxen

\$Id: LICENSE.txt,v 1.5 2006/02/05 21:49:04 elharo Exp \$

Copyright 2003-2006 The Werken Company. All Rights Reserved.Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Jaxen Project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

 THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

WSDL4I

Common Public License v1.0

03 Mar 2005

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS COMMON PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

DEFINITIONS

"Contribution" means:

in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and

in the case of each subsequent Contributor:

changes to the Program, and

additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution 'originates' from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor's behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

"Contributor" means any person or entity that distributes the Program.

"Licensed Patents" mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

"Program" means the Contributions distributed in accordance with this Agreement.

"Recipient" means anyone who receives the Program under this Agreement, including all Contributors.

GRANT OF RIGHTS

Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, world-wide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

it complies with the terms and conditions of this Agreement; and

its license agreement:

effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;

effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental and consequential damages, such as lost profits;

states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and

states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

it must be made available under this Agreement; and

a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor ("Commercial Contributor") hereby agrees to defend and indemnify every other Contributor ("Indemnified Contributor") against any losses, damages and costs (collectively "Losses") arising from claims, lawsuits and other legal actions brought by a

third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense. For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor's responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANT-ABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against a Contributor with respect to a patent applicable to software (including a cross-claim or counterclaim in a lawsuit), then any patent licenses granted by that Contributor to such Recipient under this Agreement shall terminate as of the date such litigation is filed. In addition, if Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsis-

tency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. IBM is the initial Agreement Steward. IBM may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published, Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved. This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

BSD License

The PostgreSQL JDBC driver is distributed under the BSD license, same as the server. The simplest explanation of the licensing terms is that you can do whatever you want with the product and source code as long as you don't claim you wrote it or sue us. You should give it a read though, it's only half a page. Copyright (c) 1997-2008, PostgreSQL Global Development Group All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3. Neither the name of the PostgreSQL Global Development Group nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Open Software License v. 3.0

This Open Software License (the "License") applies to any original work of authorship (the "Original

Work") whose owner (the "Licensor") has placed the following licensing notice adjacent to the copyright notice for the Original Work:

Licensed under the Open Software License version 3.0

- 1) **Grant of Copyright License.** Licensor grants You a worldwide, royalty-free, non-exclusive, sublicensable license, for the duration of the copyright, to do the following:
- a) to reproduce the Original Work in copies, either alone or as part of a collective work;
- b) to translate, adapt, alter, transform, modify, or arrange the Original Work, thereby creating derivative works ("Derivative Works") based upon the Original Work;
- c) to distribute or communicate copies of the Original Work and Derivative Works to the public, <u>with the proviso that copies of Original Work or Derivative Works that You distribute or communicate shall be licensed under this Open Software License;</u>
- d) to perform the Original Work publicly; and
- e) to display the Original Work publicly.
- 2) **Grant of Patent License.** Licensor grants You a worldwide, royalty-free, non-exclusive, sublicensable license, under patent claims owned or controlled by the Licensor that are embodied in the Original Work as furnished by the Licensor, for the duration of the patents, to make, use, sell, offer for sale, have made, and import the Original Work and Derivative Works.
- 3) **Grant of Source Code License.** The term "Source Code" means the preferred form of the Original Work for making modifications to it and all available documentation describing how to modify the Original Work. Licensor agrees to provide a machine-readable copy of the Source Code of the Original Work along with each copy of the Original Work that Licensor distributes. Licensor reserves the right to satisfy this obligation by placing a machine-readable copy of the Source Code in an information repository reasonably calculated to permit inexpensive and convenient access by You for as long as Licensor continues to distribute the Original Work.
- 4) Exclusions From License Grant. Neither the names of Licensor, nor the names of any contributors to the Original Work, nor any of their trademarks or service marks, may be used to endorse or promote products derived from this Original Work without express prior permission of the Licensor. Except as expressly stated herein, nothing in this License grants any license to Licensorië's trademarks, copyrights, patents, trade secrets or any other intellectual property. No patent license is granted to make, use, sell, offer for sale, have made, or import embodiments of any patent claims other than the licensed claims defined in Section 2. No license is granted to the trademarks of Licensor even if such marks are included in the Original Work. Nothing in this License shall be interpreted to prohibit Licensor from licensing under terms different from this License any Original Work that Licensor otherwise would have a right to license.
- 5) **External Deployment.** The term "External Deployment" means the use, distribution, or communication of the Original Work or Derivative Works in any way such that the Original Work or Derivative Works may be used by anyone other than You, whether those works are distributed or communicated to those persons or made available as an application intended for use over a network. As an express condition for the grants of license hereunder, You must treat any External Deployment by You of the Original Work or a Derivative Work as a distribution under section 1(c).
- 6) **Attribution Rights.** You must retain, in the Source Code of any Derivative Works that You create, all copyright, patent, or trademark notices from the Source Code of the Original Work, as well as any notices of licensing and any descriptive text identified therein as an "Attribution Notice." You must cause the Source Code for any Derivative Works that You create to carry a prominent Attribution Notice reasonably calculated to inform recipients that You have modified the Original Work.
- 7) Warranty of Provenance and Disclaimer of Warranty. Licensor warrants that the copyright in and to the Original Work and the patent rights granted herein by Licensor are owned by the Licensor or are sublicensed to You under the terms of this License with the permission of the contributor(s) of those copyrights and patent rights. Except as expressly stated in the immediately preceding sentence, the Original Work is provided under this License on an "AS IS" BASIS and WITHOUT WARRANTY, either express or implied, including, without limitation, the warranties of non-infringement, merchantability or fitness for

- a particular purpose. THE ENTIRE RISK AS TO THE QUALITY OF THE ORIGINAL WORK IS WITH YOU. This DISCLAIMER OF WARRANTY constitutes an essential part of this License. No license to the Original Work is granted by this License except under this disclaimer.
- 8) **Limitation of Liability.** Under no circumstances and under no legal theory, whether in tort (including negligence), contract, or otherwise, shall the Licensor be liable to anyone for any indirect, special, incidental, or consequential damages of any character arising as a result of this License or the use of the Original Work including, without limitation, damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses. This limitation of liability shall not apply to the extent applicable law prohibits such limitation.
- 9) Acceptance and Termination. If, at any time, You expressly assented to this License, that assent indicates your clear and irrevocable acceptance of this License and all of its terms and conditions. If You distribute or communicate copies of the Original Work or a Derivative Work, You must make a reasonable effort under the circumstances to obtain the express assent of recipients to the terms of this License. This License conditions your rights to undertake the activities listed in Section 1, including your right to create Derivative Works based upon the Original Work, and doing so without honoring these terms and conditions is prohibited by copyright law and international treaty. Nothing in this License is intended to affect copyright exceptions and limitations (including it fair use it fair use it fair dealing it f
- 10) **Termination for Patent Action.** This License shall terminate automatically and You may no longer exercise any of the rights granted to You by this License as of the date You commence an action, including a cross-claim or counterclaim, against Licensor or any licensee alleging that the Original Work infringes a patent. This termination provision shall not apply for an action alleging patent infringement by combinations of the Original Work with other software or hardware.
- 11) Jurisdiction, Venue and Governing Law. Any action or suit relating to this License may be brought only in the courts of a jurisdiction wherein the Licensor resides or in which Licensor conducts its primary business, and under the laws of that jurisdiction excluding its conflict-of-law provisions. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any use of the Original Work outside the scope of this License or after its termination shall be subject to the requirements and penalties of copyright or patent law in the appropriate jurisdiction. This section shall survive the termination of this License.
- 12) Attorneysi¿½ Fees. In any action to enforce the terms of this License or seeking damages relating thereto, the prevailing party shall be entitled to recover its costs and expenses, including, without limitation, reasonable attorneys' fees and costs incurred in connection with such action, including any appeal of such action. This section shall survive the termination of this License.
- 13) **Miscellaneous.** If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable.
- 14) **Definition of "You" in This License.** "You" throughout this License, whether in upper or lower case, means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, "You" includes any entity that controls, is controlled by, or is under common control with you. For purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.
- 15) **Right to Use.** You may use the Original Work in all ways not otherwise restricted or conditioned by this License or by law, and Licensor promises not to interfere with or be responsible for such uses by You.
- 16) **Modification of This License.** This License is Copyright it 2½ 2005 Lawrence Rosen. Permission is granted to copy, distribute, or communicate this License without modification. Nothing in this License permits You to modify this License as applied to the Original Work or to Derivative Works. However, You may modify the text of this License and copy, distribute or communicate your modified version (the "Modified License") and apply it to other original works of authorship subject to the following conditions:

(i) You may not indicate in any way that your Modified License is the "Open Software License" or "OSL" and you may not use those names in the name of your Modified License; (ii) You must replace the notice specified in the first paragraph above with the notice "Licensed under <insert your license name here>" or with a notice of your own that is not confusingly similar to the notice in this License; and (iii) You may not claim that your original works are open source software unless your Modified License has been approved by Open Source Initiative (OSI) and You comply with its license review and certification process.

Mozilla Public License Version 1.1

1. Definitions.

1.0.1. "Commercial Use"

means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor"

means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version"

means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code"

means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism"

means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable"

means Covered Code in any form other than Source Code.

1.6. "Initial Developer"

means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work"

means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License"

means this document.

1.8.1. "Licensable"

means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications"

means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code"

means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims"

means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code"

means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your")

means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. Source Code License.

2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof). the licenses granted in this Section 2.1 (

a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

Notwithstanding Section 2.1 (

b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor Version (or portions of such combination).

the licenses granted in Sections 2.2 (

- a) and 2.2 (b) are effective on the date Contributor first makes Commercial Use of the Covered Code. Notwithstanding Section 2.2 (
- b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party modifica-

tions of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4. Intellectual Property Matters

(a) Third Party Claims

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the **legal** file.

(c) Representations.

Contributor represents that, except as disclosed pursuant to Section 3.4 (a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient

rights to grant the rights conveyed by this License.

3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear than any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Sections 3.1, 3.2, 3.3, 3.4 and 3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the **legal** file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

6. Versions of the License.

6.1. New Versions

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.

6.3. Derivative Works

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Public License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. Disclaimer of warranty

Covered code is provided under this license on an "as is" basis, without warranty of any kind, either expressed or implied, including, without limitation, warranties that the covered code is free of defects, merchantable, fit for a particular purpose or non-infringing. The entire risk as to the quality and performance of the covered code is with you. Should any covered code prove defective in any respect, you (not the initial developer or any other contributor) assume the cost of any necessary servicing, repair or correction. This disclaimer of warranty constitutes an essential part of this license. No use of any covered code is authorized hereunder except under this disclaimer.

8. Termination

- 8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.
- 8.2. If You initiate litigation by asserting a patent infringement claim (excluding declatory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:
- such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.
- any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that Participant.
- 8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.
- 8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior

to termination shall survive termination.

9. Limitation of liability

Under no circumstances and under no legal theory, whether tort (including negligence), contract, or otherwise, shall you, the initial developer, any other contributor, or any distributor of covered code, or any supplier of any of such parties, be liable to any person for any indirect, special, incidental, or consequential damages of any character including, without limitation, damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses, even if such party shall have been informed of the possibility of such damages. This limitation of liability shall not apply to liability for death or personal injury resulting from such party's negligence to the extent applicable law prohibits such limitation. Some jurisdictions do not allow the exclusion or limitation of incidental or consequential damages, so this exclusion and limitation may not apply to you.

10. U.S. government end users

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

11. Miscellaneous

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to disputes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. Responsibility for claims

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. Multiple-licensed code

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the MPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

Exhibit A - Mozilla Public License.

"The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at http://www.mozilla.org/MPL/

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is	
----------------------	--

The Initial Developer of the Origin	nal Code is	
Portions created by	are Copyright (C)	All
Rights Reserved.		
Contributor(s):	·	
License"), in which case the provi wish to allow use of your version others to use your version of this above and replace them with the	s file may be used under the terms of theisions of [] License are applicable instead of this file only under the terms of the [] List file under the MPL, indicate your decision by detenotice and other provisions required by the [cipient may use your version of this file under either the may use your version of the may use your vers	of those above. If you cense and not to allow eleting the provisions _] License. If you do not
NOTE: The text of this Exhibit A m	nay differ slightly from the text of the notices in	the Source Code files of

the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code

GNU GENERAL PUBLIC LICENSE

Source Code for Your Modifications.

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc. 675 Mass Ave, Cambridge, MA 02139, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

O. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

- 2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

- 3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

- 4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- 5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
- 6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
- 7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and

any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

- 8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
- 9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

- 11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
- 12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix: How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.>
Copyright (C) 19yy <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) 19yy name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'. This is free software, and you are welcome to redistribute it under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program. You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program 'Gnomovision' (which makes passes at compilers) written by James Hacker.

<signature of Ty Coon>, 1 April 1989 Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

Fonts are (c) Bitstream (see below). DejaVu changes are in public domain. Glyphs imported from Arev fonts are (c) Taymjong Bah (see below)

Bitstream Vera Fonts Copyright

Copyright (c) 2003 by Bitstream, Inc. All Rights Reserved. Bitstream Vera is a trademark of Bitstream, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distribute the Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or characters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Bitstream" or the word "Vera".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been modified and is distributed under the "Bitstream Vera" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL BITSTREAM OR THE GNOME FOUNDATION BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the names of Gnome, the Gnome Foundation, and Bitstream Inc., shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from the Gnome Foundation or Bitstream Inc., respectively. For further information, contact: fonts at gnome dot org.

Arev Fonts Copyright

Copyright (c) 2006 by Tavmjong Bah. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distribute the modifications to the Bitstream Vera Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or characters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Tavmjong Bah" or the word "Arev".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been modified and is distributed under the "Tavmjong Bah Arev" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL TAVMJONG BAH BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the name of Tavmjong Bah shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from Tavmjong Bah. For further information, contact: tavmjong @ free. fr.

\$Id: LICENSE 2133 2007-11-28 02:46:28Z lechimp \$

Java Util ConcurrentPublic Domain

ICU (license below)

ICU License - ICU 1.8.1 and later

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1995-2011 International Business Machines Corporation and others

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

iTextLGPL JavaAssistLGPL JSoupApache 2.0
JMDNSApache 2.0
Password Complexity Validator(Fork of Apache 2.0 product)
Exalead, SnowTide, Oracle JDBC, NetApp, Caringo, Centera (Commercial License)
Apache POI Apache 2.0
PostGres JDBCBSD
Aperture (WP Extractor) (Special License, below)

Aperture License (for releases from 1.3.0 onward)

The Aperture Framework itself is distributed under a BSD-style license. The copyright holder for each particular source file is given within the file itself.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution

Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Third party libs

Licenses for third party libraries used in Aperture can be found in ./lib/

The use of BSD in the core framework allows developers to license implementations of Aperture APIs under any license they see fit, including proprietary and commercial licenses.

Rationale

The aperture framework is free to use and to extend. Through use of BSD, we experience that the framework is more useful and users are willing to contribute to the project.

SLF4JMIT License, below

Copyright (c) 2004-2008 QOS.ch All rights reserved. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Soft-

ware is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Spring FrameworkApache 2.0 MS SQL JDBC(Commercial License) Metro JAXWS(CDDL License, below)

COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL) Version 1.1

- 1. Definitions.
- 1.1. "Contributor" means each individual or entity that creates or contributes to the creation of Modifications.
- 1.2. "Contributor Version" means the combination of the Original Software, prior Modifications used by a Contributor (if any), and the Modifications made by that particular Contributor.
- 1.3. "Covered Software" means (a) the Original Software, or (b) Modifications, or (c) the combination of files containing Original Software with files containing Modifications, in each case including portions thereof.
 - 1.4. "Executable" means the Covered Software in any form other than Source Code.
- 1.5. "Initial Developer" means the individual or entity that first makes Original Software available under this License.
- 1.6. "Larger Work" means a work which combines Covered Software or portions thereof with code not governed by the terms of this License.
 - 1.7. "License" means this document.
- 1.8. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.
 - 1.9. "Modifications" means the Source Code and Executable form of any of the following:
- A. Any file that results from an addition to, deletion from or modification of the contents of a file containing Original Software or previous Modifications;
 - B. Any new file that contains any part of the Original Software or previous Modification; or
 - C. Any new file that is contributed or otherwise made available under the terms of this License.
 - 1.10. "Original Software" means the Source Code and Executable form of computer software code that

is originally released under this License.

- 1.11. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.
- 1.12. "Source Code" means (a) the common form of computer software code in which modifications are made and (b) associated documentation included in or with such code.
- 1.13. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. License Grants.

2.1. The Initial Developer Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, the Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license:

- (a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer, to use, reproduce, modify, display, perform, sublicense and distribute the Original Software (or portions thereof), with or without Modifications, and/or as part of a Larger Work; and
- (b) under Patent Claims infringed by the making, using or selling of Original Software, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Software (or portions thereof).
- (c) The licenses granted in Sections 2.1(a) and (b) are effective on the date Initial Developer first distributes or otherwise makes the Original Software available to a third party under the terms of this License.
- (d) Notwithstanding Section 2.1(b) above, no patent license is granted: (1) for code that You delete from the Original Software, or (2) for infringements caused by: (i) the modification of the Original Software, or (ii) the combination of the Original Software with other software or devices.

2.2. Contributor Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license:

- (a) under intellectual property rights (other than patent or trademark) Licensable by Contributor to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof), either on an unmodified basis, with other Modifications, as Covered Software and/or as part of a Larger Work; and
- (b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: (1) Modifications made by that

Contributor (or portions thereof); and (2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

- (c) The licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first distributes or otherwise makes the Modifications available to a third party.
- (d) Notwithstanding Section 2.2(b) above, no patent license is granted: (1) for any code that Contributor has deleted from the Contributor Version; (2) for infringements caused by: (i) third party modifications of Contributor Version, or (ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or (3) under Patent Claims infringed by Covered Software in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Availability of Source Code.

Any Covered Software that You distribute or otherwise make available in Executable form must also be made available in Source Code form and that Source Code form must be distributed only under the terms of this License. You must include a copy of this License with every copy of the Source Code form of the Covered Software You distribute or otherwise make available. You must inform recipients of any such Covered Software in Executable form as to how they can obtain such Covered Software in Source Code form in a reasonable manner on or through a medium customarily used for software exchange.

3.2. Modifications.

The Modifications that You create or to which You contribute are governed by the terms of this License. You represent that You believe Your Modifications are Your original creation(s) and/or You have sufficient rights to grant the rights conveyed by this License.

3.3. Required Notices.

You must include a notice in each of Your Modifications that identifies You as the Contributor of the Modification. You may not remove or alter any copyright, patent or trademark notices contained within the Covered Software, or any notices of licensing or any descriptive text giving attribution to any Contributor or the Initial Developer.

3.4. Application of Additional Terms.

You may not offer or impose any terms on any Covered Software in Source Code form that alters or restricts the applicable version of this License or the recipients' rights hereunder. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, you may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.5. Distribution of Executable Versions.

You may distribute the Executable form of the Covered Software under the terms of this License or

under the terms of a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable form does not attempt to limit or alter the recipient's rights in the Source Code form from the rights set forth in this License. If You distribute the Covered Software in Executable form under a different license, You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.6. Larger Works.

You may create a Larger Work by combining Covered Software with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Software.

4. Versions of the License.

4.1. New Versions.

Oracle is the initial license steward and may publish revised and/or new versions of this License from time to time. Each version will be given a distinguishing version number. Except as provided in Section 4.3, no one other than the license steward has the right to modify this License.

4.2. Effect of New Versions.

You may always continue to use, distribute or otherwise make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. If the Initial Developer includes a notice in the Original Software prohibiting it from being distributed or otherwise made available under any subsequent version of the License, You must distribute and make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. Otherwise, You may also choose to use, distribute or otherwise make the Covered Software available under the terms of any subsequent version of the License published by the license steward.

4.3. Modified Versions.

When You are an Initial Developer and You want to create a new license for Your Original Software, You may create and use a modified version of this License if You: (a) rename the license and remove any references to the name of the license steward (except to note that the license differs from this License); and (b) otherwise make it clear that the license contains terms which differ from this License.

5. DISCLAIMER OF WARRANTY.

COVERED SOFTWARE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED SOFTWARE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGING. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED SOFTWARE IS WITH YOU. SHOULD ANY COVERED SOFTWARE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

6. TERMINATION.

- 6.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.
- 6.2. If You assert a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You assert such claim is referred to as "Participant") alleging that the Participant Software (meaning the Contributor Version where the Participant is a Contributor or the Original Software where the Participant is the Initial Developer) directly or indirectly infringes any patent, then any and all rights granted directly or indirectly to You by such Participant, the Initial Developer (if the Initial Developer is not the Participant) and all Contributors under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively and automatically at the expiration of such 60 day notice period, unless if within such 60 day period You withdraw Your claim with respect to the Participant Software against such Participant either unilaterally or pursuant to a written agreement with Participant.
- 6.3. If You assert a patent infringement claim against Participant alleging that the Participant Software directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.
- 6.4. In the event of termination under Sections 6.1 or 6.2 above, all end user licenses that have been validly granted by You or any distributor hereunder prior to termination (excluding licenses granted to You by any distributor) shall survive termination.

7. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED SOFTWARE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

8. U.S. GOVERNMENT END USERS.

The Covered Software is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" (as that term is defined at 48 C.F.R. § 252.227-7014(a)(1)) and "commercial computer software documentation" as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Software with only those rights set forth herein. This U.S. Government Rights clause is in lieu of, and supersedes, any other FAR, DFAR, or other clause or provision that

addresses Government rights in computer software under this License.

9. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by the law of the jurisdiction specified in a notice contained within the Original Software (except to the extent applicable law, if any, provides otherwise), excluding such jurisdiction's conflict-of-law provisions. Any litigation relating to this License shall be subject to the jurisdiction of the courts located in the jurisdiction and venue specified in a notice contained within the Original Software, with the losing party responsible for costs, including, without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License. You agree that You alone are responsible for compliance with the United States export administration regulations (and the export control laws and regulation of any other countries) when You use, distribute or otherwise make available any Covered Software.

10. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

NOTICE PURSUANT TO SECTION 9 OF THE COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL)

The code released under the CDDL shall be governed by the laws of the State of California (excluding conflict-of-law provisions). Any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California and the state courts of the State of California, with venue lying in Santa Clara County, California.