

GWAVA Inc.

GWAVA Retain

For GroupWise®

Windows Install Guide

For version 1.7

Contents

IMPORTANT NOTES.....	4
Overview	7
What Retain does.....	7
How Retain works	7
Architecture	7
Best Practices for Component Placement	8
How Retain Stores the Archives.....	8
Other components that Retain depends on:	8
System Requirements	9
Supported Windows Versions.....	9
Novell GroupWise Requirements:	9
Server hardware, memory, and storage	10
2 TB Storage Limit – File System Considerations	10
Supported SQL Database Servers	10
Backup and Restore of Database and Retain.....	11
Known Incompatibilities	11
Design Considerations	12
SQL Database Server	12
Retain Server	12
Retain Worker	12
GroupWise Address Book	13
GroupWise	13
Retain-Tomcat.....	13
Installation Preparation	14
Preparing Your Server(s)	14
Java.....	14
Apache	14
Tomcat	14
SQL Server	14
Retain Installation	15
Install using Apache 2.2	15

Installing on Windows using MS IIS as the web server	17
Appendix A	18
SQL Servers	18
MySQL Quick Start Guide	18
Getting MySQL	18
MySQL Optimization	18
Installation Procedure on Windows	19
Create a new database for Retain	23
SQL Server 2005	24
SQL Server 2000	24
ORACLE 10/11	25
PostgreSQL 8.2	27
Appendix B	28
Tomcat Memory Tuning	28
Checking memory:	29
Appendix C	30
Estimating Storage Requirements	30
Appendix D	32
Uninstalling Retain	32
Appendix E	33
Some notes on running Apache 2.2 on Win32	33
Mint a CSR	34
Mint your own self-signed certificate	34
Finish off as follows	34

Technical Support

If you have a technical support question, please consult the GWAVA Technical Support section of our website at www.gwava.com.

Sales

To contact a Beginfinite 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. Beginfinite Inc. assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation. GroupWise is a registered trademark of Novell, and copyrighted by Novell. Windows is copyrighted by Microsoft. © 2005 GWAVA Inc. All rights reserved. ® GWAVA is a registered trademark.

IMPORTANT NOTES

Current Retain Customers, Retain 1.6 Upgrade

Retain 1.7 contains several important changes from Retain 1.5, and performs these updates the first time 1.7 is run. To run the update, run Retain 1.7 after the update and login as administrator. Retain will then perform the update. **DO NOT INTERRUPT THE UPDATE OR DATABASE CORRUPTION MAY RESULT.** Wait until the Retain console notifies you that the update has completed, then restart Tomcat.

Updates

Deletion Manager

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

There are two major updates to the system:

➤ **Indexer updated**

Lucene, the indexing engine that Retain uses to catalog and index the mail archive, has been updated from version 2.2 to 2.3. The update changes the index schema, providing performance enhancements.

Because of the update, once 1.7 has been installed and run, it is impossible to return to a previous version, as the previous version will view the index as broken. **The upgrade is permanent.**

➤ **Database update**

There is a database update between 1.7 and previous versions. With large databases, the update can take hours. When you first connect to the Retain Server console after the upgrade, you will be prompted to initiate the database update. This is the only screen available until the database update is complete.

DO NOT INTERRUPT THE UPDATE OR DATABASE CORRUPTION MAY RESULT. If any errors occur, Retain lists them in the Server logs. The Retain Server logs are nested with the embedded Tomcat 5 logs, which differs depending on installation platform.

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

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

If you are running your MySQL on Linux, see the important schema note in the Linux install guide, under the MySQL quick start section.

Retain Server moving to 64-bit architecture

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

The list of supported operating systems is as follows:

SuSE Enterprise Linux 9 SP4 32-bit

SuSE Enterprise Linux 10 SP2 32-bit

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

SuSE Enterprise Linux 10 SP2 64-bit

Open Enterprise Server 2 64-bit

NetWare 6.5 (Retain Worker only)

Microsoft Windows Server 2000, 2003 32-bit

Microsoft Windows Server 2003 64-bit

New Supported Software or features:

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

Retain also supports NetApp.

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

Retain also supports auto-creation of users in Livelink.

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

GroupWise 8 Support

Retain 1.7 provides full support for GroupWise 8

Java Support

The supported versions of Java have been unified for the Retain Server platform. **SUN Java 1.5, (also called java 5), is the only supported version and is required for full functionality, IBM's Java platforms are not supported.** Retain 1.7 will not work with Java 1.4x or Java 1.6. It is strongly recommended that the JDK package is installed and used instead of the JRE package in favor of the diagnostic tools included with the JDK software package.

Stubbing Server

Retain now has an integrated server to take advantage of the stubbing abilities of GroupWise 8. The Stubbing feature in GroupWise 8 allows Retain to replace archived messages and items in the GroupWise Post office that exceed a certain size with smaller 'stubs' that link back to the full message or item in the Retain Archive. The end user experience is not changed. Stubbing is not appropriate for every system. **To use stubbing, it is strongly recommended that the GroupWise system be at least 8.0HP1. Previous versions of GroupWise do not contain critical stability fixes.**

Tomcat 5.5.27 Now Embedded With Retain

Retain now ships with Tomcat embedded. Retain 1.6 requires you to install and use the embedded Tomcat. Installation on an existing Retain system running a standard Tomcat installation will overwrite the current Tomcat installation with the embedded Retain-Tomcat. Any settings or files in the current Tomcat installation directory will be removed. Installation instructions are located in the [install section](#). The listen ports for Retain Tomcat have been changed, and run as follows:

Listen ports

HTTP:	48080
AJP:	48009
Home Directory:	C:\Program Files\Apache Software Foundation\Tomcat5.5
Log Directory:	C:\Program Files\Apache Software Foundation\Tomcat5.5\logs
Service runner:	C:\Program Files\Apache Software Foundation\Tomcat5.5\bin\tomcat5.exe
Configuration file:	C:\Program Files\Apache Software Foundation\Tomcat5.5
Memory Tuning:	C:\Program Files\Apache Software Foundation\Tomcat5.5\bin\tomcat5w.exe

Memory tuning for the Retain-Tomcat has been set 256MB initial memory and 1024MB maximum.

Overview

Retain is GWAVA's long-term archiving solution and retention system designed specifically for Novell GroupWise 7.0.1 and higher. Retain uses a web-based, centralized management system to schedule and configure archive jobs. Retain can work with GroupWise retention and purge systems and settings and is designed as a cross-platform solution with full auditing built in.

What Retain does

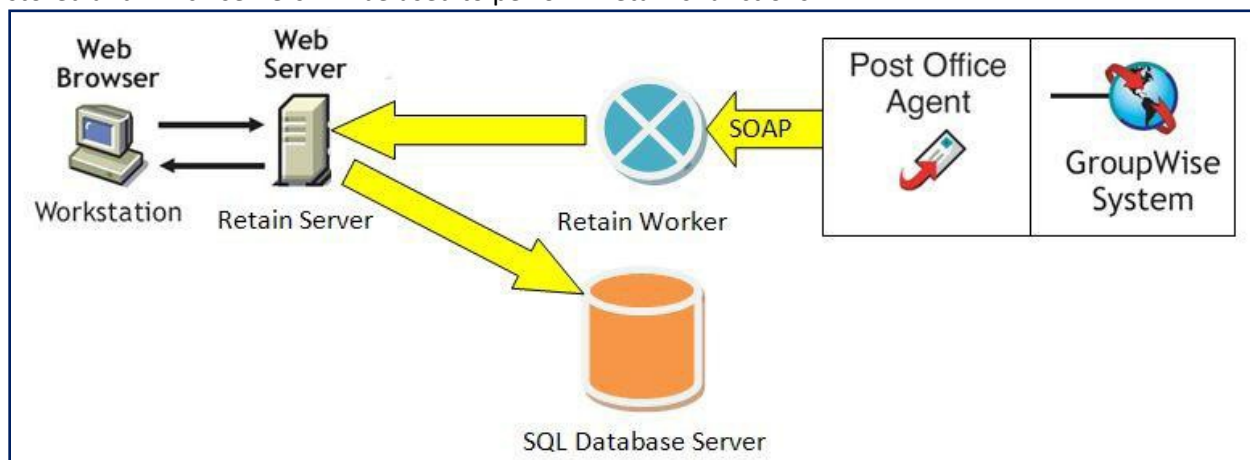
Retain provides a service of long-term storage of data as well as providing search, retrieval, and review services for retained messages. Retain is NOT a backup or emergency restoration system. (For this, we recommend GWAVA Reload.) Retain dredges messages from GroupWise post offices and stores them for long-term reference. Users may log in and review their personal archived mail and search through it.

How Retain works

The Retain Worker process uses *SOAP* and a trusted application key to collect GroupWise data as defined in a specified profile and according to a defined schedule. This data is handed to the Retain Server which in turn stores the collected data in an administrator-defined location. Users log in to the Retain Server to search through and access their archived messages. User rights are managed by the administrator.

Architecture

Retain consists of three main parts which can be installed on the same server or they can be spread out across different servers. This gives the administrator enormous flexibility in choosing where data gets stored and which servers will be used to perform Retain's functions.



- **Retain Server:** This is the heart and soul of Retain itself. All functions are controlled from the Retain Server. The server also manages the Retain Worker and stores its data in the database server. Usually there is only one Retain Server per system.
- **Retain Worker:** There is at least one per system, usually installed on each computer running a GroupWise Post Office Agent. (GWPOA) The worker does the actual data collection work. 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 itself. Retain was designed to support many different databases. The default is MySQL 5.0, but Retain also supports MySQL 5.0, Oracle 10 and 11, MS SQL server 2000, MS SQL server 2005, and PostgreSQL 8.2+

Installation, maintenance, tuning and backup of the database is the customer's responsibility.

Best Practices for Component Placement

Retain components communicate via TCP/IP. Placing all the components in the same server will yield the best communication speed. However, this is impractical in a large scale environment. Thus, you must balance the performance of the components on the physical servers with the speed of the network links joining them.

How Retain Stores the Archives

Retain uses a hybrid data storage approach. The database contains all the metadata, 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.

- Windows Server 2000 SP4+, Windows Server 2003 SP1+
- Novell GroupWise
- Apache Web Server OR Microsoft IIS
- Sun's Java 1.5 JDK
- SQL Database for storage (MySQL 5.0, Oracle 10/11, MS SQL 2000/2005, or PostgreSQL 8.2)

System Requirements

The hardest part about installing Retain is preparing GroupWise and your server(s). You must first decide *where* you want everything 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 components on different servers or run them all on one.

Minimum System Requirements

- Pentium 4 (or higher) at least 3.0GHz
- 2 GB of Ram (for the Server) 64MB of Ram (for the Worker)
- Apache 2.0x, Or IIS Version 5 or 6.
- MySQL 5.0x (or alternate supported database version)
- Sun's Java JDK 1.5x
 - Note: Java 1.6x is NOT SUPPORTED. Check to make sure it is not installed.
- As much storage space as system requires. (see [Estimating storage requirements](#))

Suggested Requirements

- Dual-Core 3 GHz or better Processor
- 4 GB of system Ram or more installed
- 160MB Ram free for Retain Worker

Supported Windows Versions

- MS Windows 2000 SP4+
- MS Windows 2003 SP1+ (32 and 64-bit)

Novell GroupWise Requirements:

- GroupWise 8, 7. (For GroupWise 7 Retain requires 7.03 or higher. 7.01 and above may function with Retain, but are no longer tested. 7.03 contains significant bug fixes and enhanced SOAP performance. 7.03HP1 recommended because it fixes a large memory leak in the original 7.03 release.)
- 7.0.3 Or later is recommended for bug fixes and performance improvements.
- SOAP port must be enabled on all POAs to be scanned for messages.
- SSL is ok, but costs significantly in performance.

Server hardware, memory, and storage

At a minimum the Retain Server component must be installed on a 2 GB RAM machine, with 1 GB RAM reserved exclusively for the Retain Server, (applies to all platforms). In addition to base memory requirements, Retain-Tomcat will require memory proportional to the amount of users on your Retain Server. Please see the section on tuning Tomcat for details.

- *Retain-Tomcat must be adjusted prior to running Retain for GroupWise. Without adjustment, Tomcat can run out of memory, causing it to crash. This will disable all web applications that depend on it. Please see the section on memory tuning for Tomcat.*

CPU should be at least a Pentium 4 or higher, at least 1.0 GHz, *the faster, the better.*

- *A considerably more powerful CPU, such as a dual core 3+ GHz, is highly recommended.*

Disk storage needs are diverse and vary widely from system to system due to mail flow and type of mail in the respective system. Retain for GroupWise compresses data as it stores, much like GroupWise does. So the storage requirements will be similar to GroupWise. Bear in mind however, that storage requirements are dependent on WHAT you decide to store and how much data you add to the GroupWise system as time moves forward. Remember that Retain for GroupWise storage is cumulative and nothing is ever lost. Generally, assuming you will store “*everything*”, the initial storage requirements will be 100% of your post office size. If you want an idea of how much more you will need as time goes on, you will have to factor in how much your post office has grown over time and deduct how much is generally pruned or deleted from GroupWise.

2 TB Storage Limit – File System Considerations

It is important to consider that some file storage systems have a limit of 2 TB. If your data storage needs will exceed this amount, please choose a file system that can accommodate this. This is something you should verify on all parts of the storage system you choose, from the hardware on up. If you believe your data storage requirements will exceed 2TB, you may wish to consider an enterprise-class storage subsystem like a SAN. Many excellent choices exist in the marketplace. Be sure your chosen platform for Retain Server can use it. This is a function of the operating system. Windows server 2003 and higher can manage more than 2 TB of storage, but this requires special partitioning to accomplish.

Supported SQL Database Servers

- MySQL 5.0x
- Oracle 10/11
- MS SQL Server 2000 / 2005
- PostgreSQL 8.2+

Both 32-bit and 64-bit versions of the above databases are supported. A 64-bit dedicated database server tuned to use most of the Server’s RAM is recommended wherever possible. As with all database servers, the more RAM the better the performance. Database Servers usually rank their priorities, in order: RAM, disk speed, and CPU speed. Some also recommend placing transaction log files on a separate volume. Such configuration is beyond the scope of this manual and the responsibility of the customer.

Backup and Restore of Database and Retain

- Backup and restore of the SQL database is your responsibility and is beyond the scope of this manual.
- Retain for GroupWise has built-in routines to backup its indexes and configuration database.
 - Your backup strategy should focus on these backup files. Doing this, you can back up the system while Retain for GroupWise is running.

Known Incompatibilities

The following are systems known to be incompatible with Retain. DO NOT USE THESE on the same server as Retain.

- Apache 1.x – Retain will NOT function with Apache 1.x
- MySQL 4.x – DO NOT USE THIS; it will cause data corruption. Use only MySQL 5.x and above.
- GroupWise WebAccess – only incompatible if using Microsoft IIS for running Retain
- Retain cannot access mailboxes with either of the following set
 - Disabled/Expired account
 - The GroupWise SOAP API cannot log into these accounts and is subject to the same limitations as standard GroupWise clients.
 - Limited client license

Novell has addressed the limited client license issue in GroupWise 7 SP3.

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.
- **The SQL server MUST NOT depend on an NFS share. NFS locking is not reliable enough for any database system and will result in database corruption. NFS shares should not be utilized in any database system.**
- Network link between the Retain Server and SQL Database Server must be speedy and reliable.
- MUST SUPPORT UTF-8 character set, and collation (MySQL 5.0x is an example of this).
- Database will have to be manually created by the administrator, and a user account assigned with rights. See the MySQL section included below. ([Appendix A](#))
- Storage requirements: Roughly equivalent to your cumulative POs, minus OFFILES and OFINDEX.
- Backup, Clustering, Tuning: See your SQL Database Server manual. This is the customer's responsibility. Tuning an SQL Database Server can make huge performance differences.

Retain Server

- Preferably installed on the same box as the POA for communications speed, reliable address book caching and authentication. This POA can be a low usage one.
 - In a larger installation, you might create a secondary domain with a dummy post office just for the purposes of speeding up the address book caching and user lookups.
 - The default POA should be highly available. A special POA just for this is a good idea.
- CPU requirements are high. The bigger and faster the better.
- Storage requirements: Approximately the same as your GroupWise Post Offices, assuming you are archiving "everything". See Appendix C for more information.
- Backup and Clustering: Backup of indexes and internal database may be done automatically. Attachments can be safely backed up manually.
- We recommend other web applications such as GroupWise WebAccess or iManager NOT be installed on this same server.
- iFolder 3 must NEVER be installed on the same box as Retain Server.

Retain Worker

- The piece that actually grabs data on a scheduled basis from the POAs is the Retain Worker. It then passes this data to Retain Server.
- Requires a reliable, speedy network connection between Retain Server and Retain Worker.
- Requires a reliable, speedy network connection to Post Office Agents being accessed.
- How many and where? Placing each worker local (on the same machine) to the SOAP POA is recommended to provide maximum network throughput and don't need SSL at the SOAP POA.

We then compress what we send to the Retain Server. So often one per POA and locally placed is better.

If you run more than one POA on the same server, you will encounter a long-standing Novell bug. You will have to manually edit the PO redirection tables. Novell doesn't recommend more than one POA on a server for a variety of reasons. You can have more than one Worker on a server but one worker will handle only one job at a time.

- To reiterate, it's recommended to put the Worker on the same server as the POA it dredges for performance reasons.

GroupWise Address Book

Retain for GroupWise gathers mail from known users. Users known to GroupWise are stored in GroupWise's *System Address Book*. Retain for GroupWise caches this information locally.

Often, the address book needs to be updated as users are added. The Retain Server component NEVER deletes a user from the cached address book.

Over time, Retain for GroupWise will know about all GroupWise users on the system, both current and past. Because Retain for GroupWise 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.

GroupWise

- SOAP MUST be enabled at all Post Office Agents being contacted by Retain Worker.
- SSL must be universally on/off at all SOAP Ports. One or the other, but not a mixed state. We can always change from plaintext to SSL, but not the reverse downgrade. However, it costs in performance. SSL in general is VERY performance intensive. Therefore we recommend leaving SSL off, and having a Retain Worker installed local to each POA in question.
- During the setup wizard, you'll need to generate a Trusted Application Key on a Win32 workstation with full rights to the primary domain. (This is in the General Admin guide)
- (optional) Retention and/or Smart Purge flags enabled.
- It's recommended that access by Retain Worker occurs at off hours, to avoid overloading the Post Office Agent.
- Address Book caching is recommended every 24 hours.

Retain-Tomcat

- Normally you can connect to the embedded Tomcat directly via port 48080. This is a good diagnostic to ensure that it is running. Make sure to allow a few minutes after the service is started before the system becomes available.
- Retain-Tomcat memory usage must be fine tuned. This cannot be stressed enough. Default values are insufficient for normal operation of Retain for GroupWise. See [tuning Tomcat section](#), below.

Installation Preparation

Preparing Your Server(s)

The general steps to preparing your server are:

- 1) Java (JDK)
- 2) Apache
- 3) Tomcat
- 4) SQL Server
- 5) You're ready to install Retain core components.

Windows

Applies to: Windows 2000 SP4+, Windows 2003 SP1+

Java

The Retain install will not complete if Java is not installed. A browser will be opened to the download site for the SUN Java 1.5 JDK. This version of Java is also called Java 5 on windows. Install normally, accepting all prompts. *We only support JDK 1.5x (Java 5) on Windows. Version 1.6 (Java 6) is NOT supported.*

Apache

Download it from apache.org (choose HTTP Server from left side). Get Apache 2.2x binary for Windows. Install normally, accepting all prompts. *We only support Apache 2.2 on Windows.*

Tomcat

The embedded Tomcat that comes with Retain installs to the standard Tomcat 5.5 directory, and therefore overwrites any configuration files you have stored there. It is recommended that the Retain Server be a dedicated machine. Retain-Tomcat may not be suitable to serve other web resources, and they are not supported. Make sure to plan your server accordingly, so any needed settings can be replicated as appropriate.

SQL Server

Retain requires a supported SQL server connection. Retain supports MySQL 5.0, Oracle 10/11, PostgreSQL 8.2+, or MS SQL Server 2000/2005.

Retain Installation

Install using Apache 2.2

***** NOTE: BE SURE YOUR SERVER'S CLOCK IS ACCURATE BEFORE YOU BEGIN! USING THE WINDOWS INTERNET TIME SERVICE IS RECOMMENDED TO KEEP YOUR CLOCK ACCURATE.** Now that you have installed Apache, Tomcat, and a SQL server, you are ready to install Retain for GroupWise's core components. Retain for GroupWise is installed as follows:

1. Verify that Apache and your SQL server are installed and working.
2. Verify that you have Java 1.5 (version 5) installed. **If the correct version of Java is not installed, the Retain installation will close and browser will open to the correct download page. Download and install the appropriate version, (32-bit, or 64-bit), for your system and restart the installer.**
3. Check the JAVA_HOME switch for your system. This is set in the Environment Variables.
Environment Variables are accessed through:
Start | Control panel | System. Select the Advanced tab and then select the 'Environment Variables' button near the bottom. JAVA_HOME should be listed in your System variables table and point to your Java 1.5 installation directory. If it is incorrect, correct the information. If it is missing, select 'New' and add the path. Syntax as follows:
Variable Name JAVA_HOME
Variable Value <Path to java installation directory> ie. C:\Program Files\Java\jdk1.5.0_13
4. Download the Retain for GroupWise package from gwava.com.
5. Unzip the package to a work directory. Almost any name will do, as long as you remember it.

Once you've downloaded the Retain Installation package and unzipped it, the files in the work directory should look like this:

ALWAYS RUN THE INSTALLER.txt	Setup Tomcat for Windows.exe
java.sh	tomcatinstall.sh
license.txt	<common>
readme.txt	<doc>
RetainInstall.sh (Linux install script)	<linux>
RetainInstall.ips (NetWare install script)	<NetWare>
RetainInstall.exe (Windows installer)	<precompile>
RetainWorker.zip (core Retain Worker)	<tools>
RetainServer.zip (core Retain Server)	<win32>
RetainStubServer.zip (core Retain Stub Server)	<win64>

6. Copy all files and folders to the server(s) you want to install to, into another work directory. This will be your "distribution" directory. After installation, you may delete this directory.
7. Run "Setup Tomcat for Windows.exe" and follow the prompts.
8. if you have not already installed or tested – Install JVM 1.5, Apache 2.2, reboot, and test them by typing http://127.0.0.1 (Apache Web Server should answer), and http://127.0.0.1:48080 (the embedded Tomcat should answer)
9. Tomcat must not be running during the Retain install. MANUALLY SHUTDOWN TOMCAT BEFORE PROCEEDING!
10. Run "RetainInstall.exe"

11. Specify the location of Apache, or IIS, and Tomcat
12. Specify where to install Retain for GroupWise (default=c:\Retain)
13. You will be prompted whether to install Retain Worker, Retain Server or both.
14. If you chose to install the Retain Server, you can now choose to install the Stub Server. **To use stubbing, it is strongly recommended that the GroupWise system be at least 8.0HP1. Previous versions of GroupWise do not contain critical stability fixes.**
15. Choose Apache 2.2x.
16. Specify the Apache 2.2x home directory
17. Specify the Tomcat directory and connection port (leave as default unless you know otherwise).
The product will install.
18. If you installed Retain Server on this machine, you will be asked if you wish to Pre-compile the JSPs. Pre-compiling JSPs can vastly improve the user interface performance and is recommended for a better experience. It may take several minutes to complete. Pre-compiling JSPs will fail, (with "cannot find compiler" message), if you have not set up JAVA_HOME correctly in step 3.

Retain-Tomcat is run as a System Service

You are now ready to go on to the general administration and use guide.

Installing on Windows using MS IIS as the web server

1. Install MySQL and set up the database.
2. After installing JDK 1.5, IIS, reboot, and test them by typing `http://127.0.0.1` (IIS Web Server should answer). **If the correct version of Java is not installed, the Retain installation will close and browser will open to the correct download page. Download and install the appropriate version, (32-bit, or 64-bit), for your system and restart the installer.**
3. Check the JAVA_HOME switch for your system. This is set in the Environment Variables.
Environment Variables are accessed through:
Start | Control panel | System. Select the Advanced tab and then select the 'Environment Variables' button near the bottom. JAVA_HOME should be listed in your System variables table and point to your Java 1.5 installation directory. If it is incorrect, correct the information. If it is missing, select 'New' and add the path. Syntax as follows:
Variable Name JAVA_HOME
Variable Value <Path to java installation directory> ie. C:\Program Files\Java\jdk1.5.0_13
4. Run "Setup Tomcat for Windows.exe" and follow the prompts. Choose IIS as the server type. (To check Retain-Tomcat, type <http://127.0.0.1:48080> Retain-Tomcat should answer.)
5. YOU MUST MANUALLY SHUTDOWN TOMCAT BEFORE PROCEEDING! (See [start scripts](#), and use the tomcat5w.exe to start and stop Retain-Tomcat, or shut down the system service through the Windows Services window. The Retain-Tomcat service is called "Apache Tomcat tomcat5".)
6. Run RetainInstall.exe
7. Choose IIS 5/6 when asked for the web server
8. Specify the location of Tomcat
9. Specify where to install Retain for GroupWise (default=C:\Retain)
10. You will be prompted whether to install Retain Worker, Retain Server or both.
The product will install.
11. If you chose to install the Retain Server, you can now choose to install the Stub Server. **To use stubbing, it is strongly recommended that the GroupWise system be at least 8.0HP1. Previous versions of GroupWise do not contain critical stability fixes.**
12. Specify the Tomcat directory and connection port (leave as default unless you know otherwise).
13. The product will install.
14. If you installed Retain Server on this machine, you will be asked if you wish to Pre-compile the JSPs. Pre-compiling JSPs can vastly improve the user interface performance and is recommended for a better experience. It may take several minutes to complete. Pre-compiling JSPs will fail, (with "cannot find compiler" message), if you have not set up JAVA_HOME correctly.
15. Reboot.

Retain-Tomcat is run as a System Service

Appendix A

SQL Servers

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 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 or 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

- 1) Go to <http://www.mysql.com>
- 2) Get MySQL 5.0.x "Essentials" Package

MySQL Optimization

MySQL requires careful memory tuning. By default it assumes it is running on a 64 MB Server. In fact, you 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.

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

<http://www.mysqlperformanceblog.com/files/presentations/UC2007-Innodb-Performance-Optimization.pdf> covers InnoDB specifically. Clear and concise; the "bible" for optimization:

<http://www.mysqlperformanceblog.com/files/presentations/UC2007-MySQL-Server-Settings-Tuning.pdf>

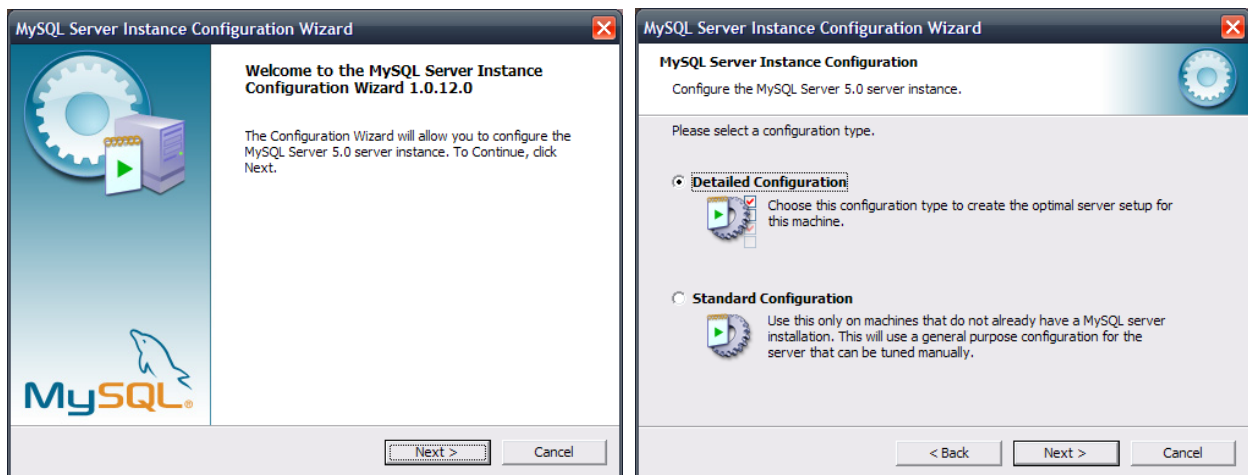

Also, for basic optimization directly after MySQL installation, see:

<http://www.mysqlperformanceblog.com/2006/09/29/what-to-tune-in-mysql-server-after-installation/>

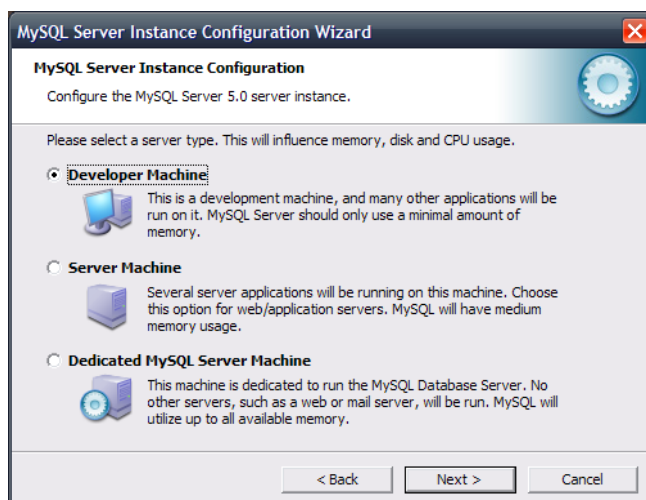
Installation Procedure on Windows

MYSQL is provided in a standard Windows Installer package.


- 1) Run the installer package.
- 2) Accept all defaults.
- 3) Go to your program list and select the MySQL Instance Configuration Wizard.



Start the Instance Configuration Wizard, and select the Detailed Configuration option.


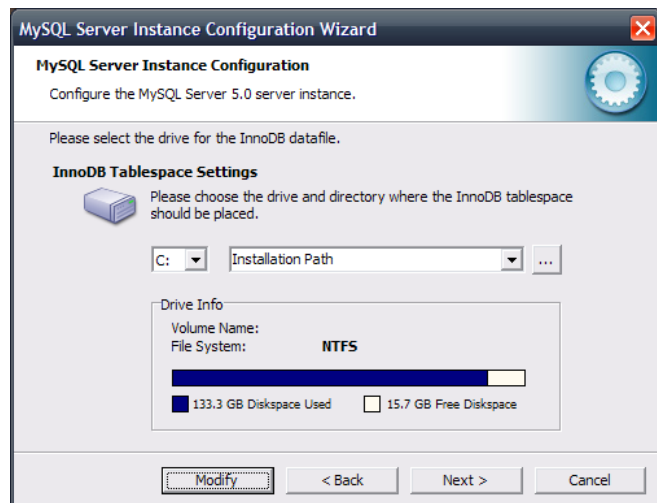


Select the appropriate configuration for the machine that is running the MySQL instance. It is recommended that you select either the Server Machine or Dedicated MySQL Server Machine setting to provide enough memory for MySQL. This will perform basic memory and CPU optimizations for your system.

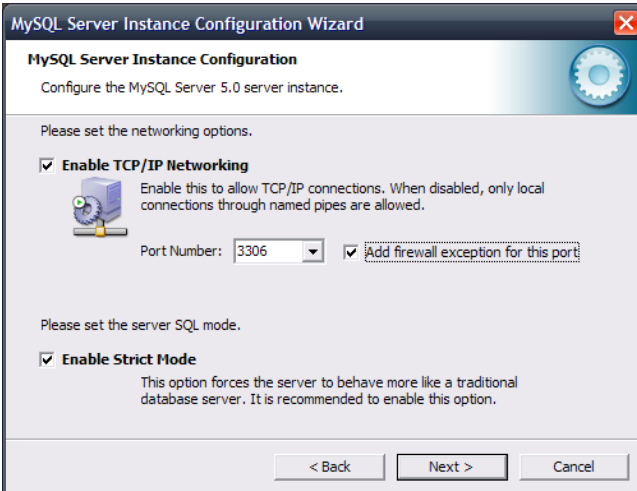


Select the database type for this instance. (If you are using this instance of MySQL for more than one database, configure accordingly. Retain uses the InnoDB storage engine.)

Chose the storage location for the Retain DB. The location you specify here will need to be large enough for your system.

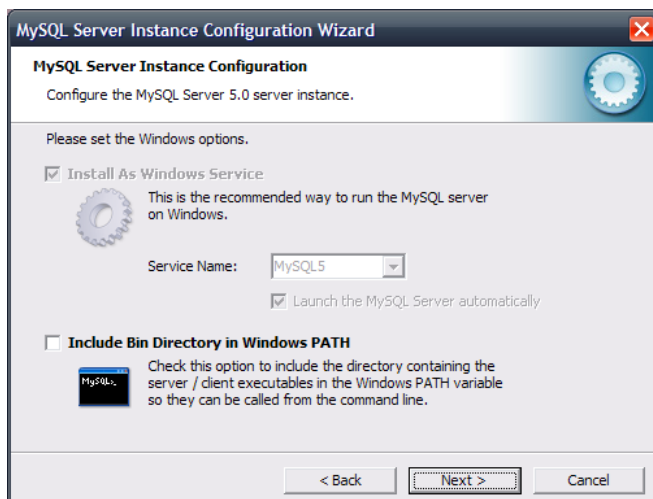
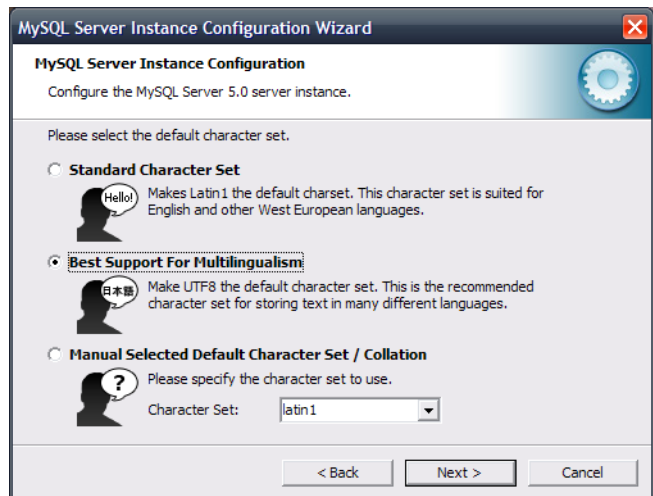


Set the desired amount of connections. (To calculate how many connections you will need, keep in mind that Retain will take more connections for multiple jobs running at the same time, plus one connection for every user logged into the archive viewer. 20 connections should be fine for Retain, plus how many users you will have searching the database at any given time. 100 connections should be sufficient for all but the largest systems.)

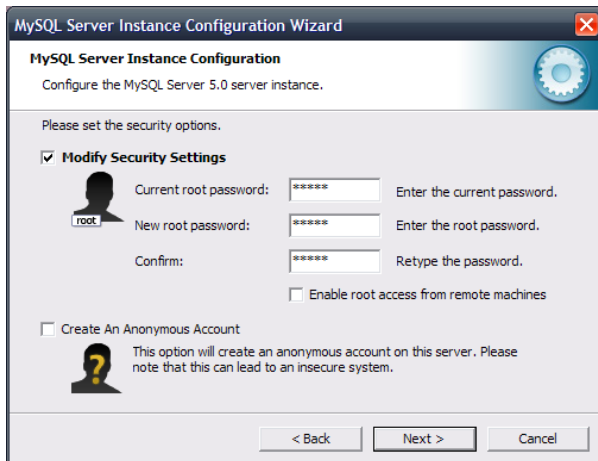


Enable the TCP/IP networking for the database and open the firewall for the port specified. Enable Strict Mode. (Take note of the port number, Retain will need to know this number for configuration performed in the general admin guide.)

Choose utf-8 for the character set.



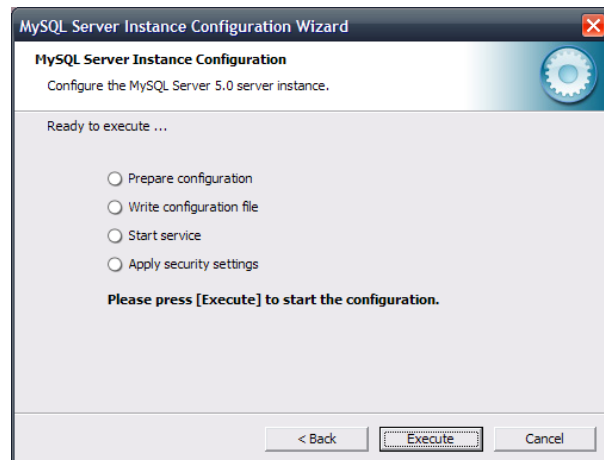
Install MySQL as a Service and select Include the Bin directory in the Windows PATH.




Enter the root passwords as prompted and configure remote root access as required.

If you have incorrectly configured any part of the installation, this is the last chance to change the settings before they are written. Select Execute to perform the listed tasks.

This will start MySQL as a service and allow you to connect using the client provided by the installation. You can then create your database.



Start the Command Line Client and login as root.



Create a new database for Retain

Applies to: All platforms

1) Log in to MySQL, as in the last step of the installation procedures.

2) Type:

```
CREATE DATABASE retain DEFAULT CHARACTER SET 'utf8' DEFAULT COLLATE 'utf8_general_ci'; *** 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, type:

```
DROP DATABASE retain; and repeat step 2.
```

Create a user for Retain and grant rights.

It is poor security practice to use the root account for ordinary database access. Retain needs a user account created that has full rights to the retain database. The example below will use mjb as the username and the password will be turnip. You may substitute these values for the user name and password you desire.

1) While logged in to MySQL, type:

```
GRANT ALL PRIVILEGES ON retain.* TO 'mjb'@ '%' IDENTIFIED BY 'turnip'; <hit enter>
```

```
GRANT ALL PRIVILEGES ON retain.* TO 'mjb'@ 'localhost' IDENTIFIED BY 'turnip'; <hit enter>
```

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

Further tuning

You need to know:

- MySQL server IP Address/HostName
- data port
- database username
- database password

It is your responsibility to learn more about your SQL Database Server. Tuning these for performance often requires an experienced DBA. It is your responsibility to backup the database store, and make it fault tolerant.

SQL Server 2005

(SP1+ required)

1. Install as default instance, with Latin 1 as standard encoding.
2. Enable SQL Server Authentication in addition to Windows Authentication
3. Verify server is listening on port 1433 (telnet ip address 1433)
4. Run SQL Server Management Studio, connect to Server
5. Under Security/Logins, right click and create a new user, (for example retainuser) that uses SQL Server authentication. Assign a password.
6. 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)

SQL Server 2000

(SP4 required)

1. Install as default instance, with Latin 1 as standard encoding.
2. Enable SQL Server Authentication in addition to Windows Authentication
3. Make sure TCPIP is enabled, on port 1433, during setup.
4. Verify server is listening on port 1433 (telnet ip address 1433)
5. Run Enterprise Manager, connect to Server
6. Under Security/Logins, right click and create a new user, (for example retainuser) that uses SQL Server authentication. Assign a password.
7. Right click on Databases, and create a new database named retain
8. Expand retain, and under users create a new user (retainuser) and assign all rights to the db
9. Right click on retain, and make sure retainuser has all of the Permissions set.

ORACLE 10/11

We support the usage of Oracle 10 and 11 on all platforms. However, we do not support the installation or tuning of Oracle for this purpose.

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\oradata\retain\temp01.dbf' SIZE 20M REUSE AUTOEXTEND ON NEXT 640K
MAXSIZE UNLIMITED
SMALLFILE UNDO TABLESPACE "UNDOTBS1" DATAFILE
'C:\oracle\product\10.2.0\oradata\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).

Tuning

Tuning Oracle is quite far beyond the scope of this manual. The manuals on Oracle's site, as well as several days of training are strongly recommended.

PostgreSQL 8.2

1. When installing, specify the UTF-8 encoding, and, if the Retain Server is on another machine, that you want to accept connections on all addresses, not just localhost
2. Connect to PostgreSQL using psql command line with the PostgreSQL user you created.
3. Create a new user. In the example below, we are creating a user called retainuser with a password of turnip

```
CREATE ROLE retainuser LOGIN ENCRYPTED PASSWORD 'password';
```

4. Create the retain database, and assign ownership to the user created in previous step.

```
CREATE DATABASE retain WITH ENCODING='UTF8' OWNER=retainuser;
```

5. If the Retain Server is on another machine, you'll also need to edit the pg_hba.conf, and restart PostgreSQL.

Here's an example entry, delimited with tabs, for a user mjb being granted rights to access retain from 192.168.2.2

```
host    retain  mjb    192.168.2.2/32 md5
```

Here is another example, the same except allowing anyone on 192.168.2.x to connect

```
host    retain  all    192.168.2.2/24 md5
```

Appendix B

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

A minimum of 1 GB of RAM is recommended. For larger systems, 2-16 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 other 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 64MB RAM is recommended, and 160MB is strongly recommended, to give room for future functionality. Remember, if this RAM is in use by Tomcat, it is unavailable to other server processes – leave sufficient physical RAM for these. Also, remember this memory is shared with other web applications.

Tomcat Configuration Parameters

There are three significant parameters:

1) Memory allocated upon Tomcat startup

- Indicates how much memory is immediately allocated and reserved to Tomcat upon startup. This memory will be in usage for the entire lifetime of Tomcat and never available to the other server processes.
- Typically this is set to 50%-100% of the maximum memory parameter discussed below. It can be lower, but pre-allocating a sizeable percentage of memory *enhances performance* and reduces memory fragmentation.

2) Maximum Memory available to Tomcat

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

3) Stack Size

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

NOTE: 32-bit Linux can allocate a maximum of 3 GB of memory for Java and Tomcat.

Retain should be tuned to use between 2 GB to 16GB of memory.

Tomcat-Retain is pre-tuned to use 256 MB initial memory, with a max of 1024 MB. This will not be enough for the server.

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

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.

Windows Tomcat memory tuning:

- 1) Go to the System Tray icon for Tomcat. (If using the embedded Retain Tomcat, launch tomcat5w.exe located in C:\Program Files\Apache Software Foundation\Tomcat 5.5\bin. Skip step 2.)
- 2) Choose configure.
- 3) Go to the Java tab.
- 4) Set the Minimum Memory (always used), Maximum Memory, and Stack Size.
(Retain Tomcat is pre-tuned to 256MB initial and 1024MB max memory with a 160K stack size.)

NOTE: 32-bit Windows can allocate a maximum of 2gb of memory to Java and Tomcat.

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.

Appendix C

Estimating Storage Requirements

How much storage space you will need in your Retain system is dependent on many factors. While there is no hard-and-fast rule, we have provided the following guide to help you evaluate your own needs. If you have GWAVA Redline, you could use its reports to help you with your estimates. In all cases, you are making an estimate for how much disk space to provide to Retain. When rounding your figures, round up to provide more disk space than you estimate you need. Calculation guide: Use the table below to help you calculate your storage requirements for Retain. Do this for each post office you plan to store in Retain and add up the results to get a final figure.

- 1) Get the disk space used by the post office, usually measurable in Gigabytes (GB). If you use GroupWise Document Management and its data is stored within the post office, run a library report using GWCHECK and subtract these values. Alternately, get the size of your post office directory and subtract the size of the DMS directory within it.
- 2) Estimate how much your post offices will grow in size over the upcoming years. There are many ways to estimate this. See below for suggestions. Remember that Retain is cumulative.
- 3) Decide how long you want to keep your data (retention time) in years.
- 4) Calculate your disk space requirements by multiplying your estimated growth per year by the number of years and adding the current size of your post office.

Parameter	Answer	Value #
Post Office Space Used (GB): (not including document management)		1
Yearly Additional Disk Space (GB): (estimate)		2
How long will you keep your data (Years)		3
Disk Space Requirements (GB): Calculate: [value 2] x [value 3] + [value 1]		4
Add 35% for indexes to get a Grand Total (GB) Calculate [value 4] x 1.35		5

Remember, this is only an ESTIMATE, assuming you will be storing EVERYTHING. As you restrict what you will actually store, this figure will go down. For example, do you want to store 50MB attachments? It is your choice.

Suggestions on methods to estimate your post office growth per year

Post office growth varies wildly from company to company and even from post office to post office. There is no hard-and-fast rule which works for everybody. However, some historical data from your own company can help you make a reasonable estimate to determine your ongoing disk space requirements. It is very strongly recommended that the file system you use to store your Retain data be either LVM, NSS or some other system which will allow you to simply add to the storage pool as your needs dictate. Retain tends to need less disk space than your GroupWise post office. It's a small amount less but a conservative estimate would be to simply say "the same amount". Naturally, this only applies if you are storing "everything". Not every organization will store "everything". If, for instance, you're only storing data for certain specific users, you can run GWCheck to determine approximately how much of the post office is used by their data. Redline has a report on post office usage which will tell you right away how much space these users are taking. Take this figure as your baseline "post office size today". Everything else will be calculated based on this figure. $\text{Current Size} = \text{Past Size} + \text{Additions} - \text{Deletions}$ Retain Needs = Past Size + Additions So the real challenge is making a reasonable estimate of what was or will be deleted. Calculate these figures for each post office you plan to store in your Retain archive and add the results together to get a final total. Suggested Method

- 1) Take your post office size today. (Present Size)
- 2) Take the size of the post office at some point in the past – as close to a year as you can get. (Past Size)
- 3) Calculate the size difference between them (Size Increase) and the time difference (Time Difference)
- 4) Extrapolate the size difference over the year. You'll get an estimate of NET growth. (mail was deleted too).
- 5) Estimate how much mail was deleted from the system. This will be a wild guess at best but only you can make such an estimate.
- 6) Your annual growth will be Net Growth + Deletions.

Even if you can measure your actual growth from last year to this year, it is a NET growth. In addition, rates of growth tend to go up as bandwidth increases because people are simply sending larger and larger items. Some people store things in GroupWise like it was a filing cabinet. So it's a good idea to overestimate your needs and plan accordingly.

Appendix D

Uninstalling Retain

Should you wish to uninstall Retain, use the following general procedure. Actual steps will vary according to your host Operating System.

1. Shut Down Tomcat
2. Shut Down Apache
3. Remove the data storage path (delete the files)
 - a. Can be: sys:retain, \retain, /var/opt/beginfinite/retain ... etc.
 - b. See the Retain Administration Guide section on storage to see where your data actually is placed.
4. Remove the program files
 - a. Stored under storage path
5. Remove the RetainWorker.xml and RetainServer.xml from Tomcat.
 - a. Windows tomcat 5.5 <installPath>\conf\Catalina\localhost
6. Remove the directories under tomcat/work (but NOT work itself)
7. Remove references to Retain for GroupWise in httpd.conf
 - a. Can be: /etc/apache2, sys:\apache2/conf, or <InstallPath>\conf).
8. (optional) remove MySQL, Apache, Tomcat
9. If Apache and Tomcat were not removed, restart them both to verify that they load without problems.

Appendix E

Some notes on running Apache 2.2 on Win32

The bugs discussed below were found in Apache 2.2.4. Since then Apache 2.2.6 has been released and may address these issues.

1. You must download the SSL enabled version.

To do so you need to:

- a) go to the binaries directory
- b) <http://www.eng.lsu.edu/mirrors/apache/httpd/binaries/win32/>
- c) download the .msi installer that has "openssl" in the name.

2. To enable SSL,

- a) Install normally
- b) shutdown the services
- c) Open the httpd.conf file in Apache2.2/conf
- d) uncomment (remove #) on these lines:

```
LoadModule ssl_module modules/mod_ssl.so
Include conf/extra/httpd-ssl.conf
```

- e) Save.

3. Open the Apache2.2/conf/extra/httpd-ssl.conf file
4. Correct all of these lines to be "quoted" or they will fail if there are spaces in the path (this is a bug in the apache installer).

Below assumes you installed Apache 2.2 to default location. Each "paragraph" is one line. It's shown on multiple lines for space reasons but a space separates one line from the other.

Note the quotation marks.

```
SSLSessionCache "shmcb:C:/Program Files/Apache Software
Foundation/Apache2.2/logs/ssl_scache(512000)"
```

```
DocumentRoot "C:/Program Files/Apache Software
Foundation/Apache2.2/htdocs"
```

```
ErrorLog "C:/Program Files/Apache Software
Foundation/Apache2.2/logs/error_log"
```

```
TransferLog "C:/Program Files/Apache Software
Foundation/Apache2.2/logs/access_log"
```

```
SSLCertificateFile "C:/Program Files/Apache Software
Foundation/Apache2.2/conf/server.crt"
```

```
SSLCertificateKeyFile "C:/Program Files/Apache Software
Foundation/Apache2.2/conf/server.key"
```

```
CustomLog "C:/Program Files/Apache Software
Foundation/Apache2.2/logs/ssl_request_log" \ "%t %h
%{SSL_PROTOCOL}x %{SSL_CIPHER}x \"%r\" %b"
```

5. The SSLCertificateFile and SSLCertificateKeyFile directives need to point to a valid certificate and key. You can either mint your own, or submit one via a CSR.

Note: You may get an error about openssl unable to find its conf file. Its conf file is in the Apache2.2/conf file and you can specify the location with -conf command line.

6. Create a directory c:\ssl, and cd to Apache2.2/bin directory in a DOS window.
7. Generate a server key

```
openssl genrsa -des3 -out server.key 1024
```

(Do NOT ever lose the password)!

8. Generate an unsecure server key (needed for Win32 apache)

```
move server.key server.key.secure  
openssl rsa -in server.key.secure -out server.key
```

9. Now, either mint a CSR (if you want to submit for to a CA for signing), or sign your own. Choose one.

Mint a CSR

```
openssl req -new -key server.key -out server.csr
```

Fill out ALL fields. The most important is the name field, which should match the server DNS (www.yourcompany.com)

Then submit the CSR to your Certificate Authority.

You'll get back another file from them, server.crt

Mint your own self-signed certificate

```
openssl req -new -x509 -nodes -sha1 -days 365 -key server.key -out server.crt
```

Finish off as follows

```
move server.* to c:\ssl
```

Now backup c:\ssl

Copy the server.key, server.crt to Apache2.2/conf