

GWAVA Inc.

GWAVA® Retain™

Windows Install Guide

Version 3

“GWAVA” is a registered trade mark of GWAVA Inc, 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1
Retain” is a trade mark of GWAVA Inc, 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1
GroupWise is a trademark of Novell, Inc.
Exchange and Windows are trademarks of Microsoft Inc.

Contents

Technical Support	4
Sales	4
Intended Audience.....	4
Organization.....	4
Copyright Notices.....	4
System Requirements	6
Supported Windows Versions.....	7
Novell GroupWise Requirements:	7
Supported SQL Database Servers	7
Known Incompatibilities	8
What Retain does.....	9
How Retain works	9
Architecture	9
Best Practices for Component Placement	10
How Retain Stores the Archives.....	10
Other components that Retain depends on	10
Design Considerations	10
SQL Database Server	10
Retain Server.....	11
Retain Worker	11
Messaging system Address Book.....	11
GroupWise	12
Embedded Tomcat	12
Installation Preparation	13
Preparing Your Server(s)	13
Java.....	13
Apache or IIS	13
Tomcat	13
SQL Server	13
Retain Installation	14

Install using Apache 2.2	14
Install using MS IIS	16
Appendix A: SQL Servers	17
MySQL Quick Start Guide.....	17
Getting MySQL	17
MySQL Optimization	17
Installation Procedure on Windows	18
Create a new database for Retain	22
SQL Server 2008 R2	23
SQL Server 2005	23
ORACLE 10/11g	24
PostgreSQL 8.2-8.4.....	26
Appendix B: Tomcat Memory Tuning	27
Checking memory:	28
Appendix C: Estimating Storage Requirements	29
Appendix D: Uninstalling Retain	30
Appendix E: Notes on running Apache 2.2 on Win32.....	30
Mint a CSR.....	31
Mint your own self-signed certificate	31
Finish off as follows.....	31
Appendix F: Windows Server 2008 and 2008 R2; IIS 7, 7.5	32

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 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 will not interfere with the standard installation and ports; they have been changed, as follows:

HTTP port:	48080
AJP port:	48009
Home Directory:	C:\Program Files\Apache Software Foundation\Tomcat
Log Directory:	C:\Program Files\Apache Software Foundation\Tomcat\logs
Start / Stop:	Tomcat service start/stop option from the Tomcat configuration window
Configuration:	Start All Programs Tomcat 5.5 Configure Tomcat
Memory Tuning:	Defaults to initial memory of 256MB and max of 1024MB. (Memory should still be tuned. See the tomcat memory tuning section in the appendix.)

The Retain install script, RetainInstall.exe, detects whether the embedded Retain-Tomcat installation is present or not, and, if missing, begins the embedded Tomcat install.

The embedded Tomcat install shuts down and removes the Retain files from the standard Tomcat installation. The embedded Tomcat is then installed. The standard Tomcat installation will still automatically start on boot if it was set to do so, and the memory that was set aside for that service will still be unavailable to the new installation. Because of the considerable memory resources that the standard Tomcat will consume, it is **STRONGLY** recommended that the memory tuning previously performed be removed, and if possible, the standard Tomcat installation completely removed or disabled. It is also recommended that the Retain Server is run on its own dedicated machine.

To ensure compatibility and performance for Retain in the future, the embedded Tomcat will be the only supported method of running Retain. Due to these tweaks, the embedded Tomcat is not suitable to run any other web services and may be incompatible.

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 components on different servers or run them all on one. VMWare ESX and ESXi 4 and 5 are fully supported, but all 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
- 256 MB Ram free for Retain Worker
- Apache 2.2
- Supported SQL Database ([Listed below](#))
- Sun Java 1.6 JDK 64-bit (32-bit platform supported for the Retain Worker only.)
- As much storage space as messaging system requires*. (see [Estimating storage requirements](#))

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**)

- 4-8 GB RAM
- Use of Exalead requires additional 4-8 GB RAM

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

500-5,000 active mailboxes:

- 4-16 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

- 4-16 GB RAM for Retain Server
- 8-64 GB RAM for Index Engine (Exalead)
- Database on appropriate server
- Retain Worker s installed remotely on each post office in messaging 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.

*The Advanced Indexing engine, Exalead, requires EXT3 or NTFS partitions

Supported Windows Versions

Server

- MS Windows 2003 SP1+ 64-bit
- MS Windows 2008 (R2) 64-bit

Worker

- MS Windows 2003 SP1+ 64-bit and SP1+ 32-bit
- MS Windows 2008 R2 64-bit and 32-bit
 - *Windows 32-bit versions supported for workers only

Microsoft Exchange Requirements:

- Exchange server 2007 SP 3
- Exchange server 2010 SP 2 or higher

Novell GroupWise Requirements:

- GroupWise 8, 7.0.3 HP 1 or above. (GroupWise 7.x should work with Retain, but anything below 7.0.4 is untested, and 7.0.3 HP 1 is **strongly** recommended for bug fixes and performance.)
- Retain Stubbing server requires GroupWise 8.01 and above.
- SOAP port must be enabled on all POAs to be scanned for messages.
- SSL is supported, but significantly impacts performance.

Supported SQL Database Servers

- MySQL 5.1x, 5.5x
- Oracle 10/11g
- MS SQL Server 2005, 2008, 2008 R2
- PostgreSQL 9.x

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.
- **iFolder 3.x** – iFolder 3.x and Retain must NOT be used on the same server. Disruption to the Retain Worker will result.
- **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 SOAP API cannot log into these accounts, both for GroupWise and Exchange, and Retain is subject to the same limitations as standard clients.
 - Limited client license

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

If installing Retain to an IIS server, please refer to the [IIS installation guide in the appendix](#).

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 dredges messages from the messaging system 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 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.

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

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 Worker s 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 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.

- 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 1.6 JDK (This is installed automatically in the installation script on systems with internet access. The JDK is recommended, as it provides additional troubleshooting options. **This is the ONLY supported version of Java.**)

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 5.5x is an example of this).

- The Retain database will have to be manually created by the administrator, and a user account must be assigned with full rights. See the Database section. ([Appendix A](#))
- Storage requirements: Roughly equivalent to the cumulative size of the message data store.
- **Backup, Clustering, Tuning is the customer's responsibility.** For specific instructions see your SQL Database Server manual. Tuning an SQL Database Server can make huge performance differences. (Between 4 and 32 GB ram should be allocated for the SQL server depending on message data requirements.)

Retain Server

- CPU requirements are high. The bigger and faster the better.
- Storage requirements: Approximately the one and a half times the size of the existing messaging system data store, assuming Retain will be archiving "everything". For Exchange size information and further GroupWise help, [see Appendix C](#).
- Backup and Clustering: Backup of Lucene indexes is performed 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 messaging system 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 CAS and Mailbox Servers or Post Office Agents being accessed.
- To provide maximum network throughput, do not use SSL at the SOAP POA or Mailbox Server. The Retain Worker then compresses the data and sends it to the Retain Server.
- DO NOT put iFolder 3 on a Retain Worker machine.
- If desired, one Retain Worker can be placed on the same box as the Server for communication performance reasons. (More than one Retain Worker sharing the same physical box as the Server is unsupported.)

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.

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 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.

All message system setup procedures are covered in the administration guide.

GroupWise

- SOAP MUST be enabled at all GroupWise 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. Therefore we recommend leaving SSL off.
- A **Trusted Application Key**, either created manually for GW 8 or later, or using the key generator from a Win32 workstation with full rights to the primary domain. (This is in the Administration Guide)
- (optional) Retention and/or Smart Purge flags enabled.
- It is recommended to run retention jobs during off hours to avoid overloading the messaging system.
- Address Book caching is recommended every 24 hours. (default)

Exchange

- Journaling must be enabled to achieve compliance.
- A mailbox enabled user with impersonation and delegation rights must be created for Retain.
- Basic authentication must be enabled in IIS on all CAS servers in the Exchange system.
- It is recommended to run retention jobs during off hours to avoid overloading the messaging system.
- Address Book caching is recommended every 24 hours. (default)

Social Messaging

- RSM Gateway must be installed and configured according to Retain for Social Media guide.
- Users must be configured so their HTTP traffic passes through RSM Gateway

Embedded Tomcat

- 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.
- **Embedded Tomcat memory usage must be fine-tuned.** This cannot be stressed enough. Default values are insufficient for normal operation of Retain. See [tuning Tomcat section](#), below.

Installation Preparation

Preparing Your Server(s)

The general steps to preparing your server are:

- 1) Java (JDK)
- 2) Apache or IIS ([See appendix](#) for IIS)
- 3) Tomcat
- 4) SQL Server

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 1.6 JDK. This version of Java is also called Java 6 on windows. Install normally, accepting all prompts.

Apache or IIS

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.*
For IIS, use the version appropriate for the platform.

Tomcat

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

SQL Server

Retain requires a supported SQL server connection. The mail storage and configuration database for Retain is housed by an SQL server. See the [courtesy SQL creation guide](#) in the appendix and the [software requirements page](#) for supported SQL servers and general database creation instructions.

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's core components. Retain is installed as follows:

1. On the SQL database server of your choice and create the retain database. See [SQL Servers](#) in the Appendix for instructions.
2. Verify that you have Apache and Java 1.6 (version 5) installed. **If the correct version of Java is not installed, the Retain installation will close and browser will download the appropriate version, (32-bit for worker only, or 64-bit), for your system. Install the downloaded Java version 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 1.6 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\jdk1.6
4. Unzip the downloaded Retain install 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 working directory should look like this:

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

5. Copy all files and folders into another work directory on the server(s) you want to install to. This will be your "distribution" directory. After installation, you may delete this directory if desired.
6. Run "Setup Tomcat for Windows.exe" and follow the prompts.

7. If you have not already installed or tested – Install JVM 1.6, 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)
8. Tomcat **must not** be running during the Retain install. MANUALLY SHUTDOWN TOMCAT BEFORE PROCEEDING!
9. Run “RetainInstall.exe”
10. Specify where to install Retain (default=C: \Retain)
11. You will be prompted whether to install Retain Worker, Retain Server or both.
12. If you chose to install the Retain Server, you can now choose to install the Stub Server. **Stubbing requires GroupWise 8.0HP1 or later. Previous versions of GroupWise do not contain critical fixes.**
13. Choose Apache 2.2x.
14. Specify the home directory of Apache.
15. Specify the Tomcat connection port (leave as default unless you know otherwise).
16. Confirm settings and select 'install'.
The product will install.
17. If Microsoft C++ 2005 Runtime distribution is not installed, the appropriate version will automatically be installed.
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.

Tomcat is run as a System Service and shows up as ‘tomcat5’ in the system services list. (Start | Run | services.msc)

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

Install using MS IIS

1. On the SQL database server of your choice and create the Retain storage database. See [SQL Servers](#) in the Appendix for instructions. (Ensure IIS is installed with ISAPI extensions and filters, covered in [Appendix F](#).)
2. After installing JDK 6, (<http://download.gwava.com/jvm/jdk-6u26-windows-x64.exe>), 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 6 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\jdk1.6.0_26
4. Run "Setup Tomcat for Windows.exe" and follow the prompts. (To check Tomcat, type <http://127.0.0.1:48080> Tomcat should answer.)
5. YOU MUST MANUALLY SHUTDOWN TOMCAT BEFORE PROCEEDING! (Use the tomcat5w.exe to start and stop Tomcat, or shut down the system service through the Windows Services window. The Tomcat service is called "Apache Tomcat tomcat5".)
6. Run "RetainInstall.exe"
7. Specify where to install Retain (default=C: \Retain)
8. You will be prompted whether to install Retain Worker, Retain Server or both.
9. If you chose to install the Retain Server, you can now choose to install the Stub Server. **Stubbing requires GroupWise 8.0HP1. Previous versions of GroupWise do not contain critical fixes.**
10. Choose IIS when asked for the web server
11. Specify the location of Tomcat
12. Specify the Tomcat directory and connection port (leave as default unless you know otherwise).
13. Confirm settings and select 'install'.
The product will install.
14. If Microsoft C++ 2005 Runtime distribution is not installed, the appropriate version will automatically be installed.
15. 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.
16. Reboot.

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 5.5 (download appropriate MSI installer for the host system, 32-bit or 64-bit.)

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.

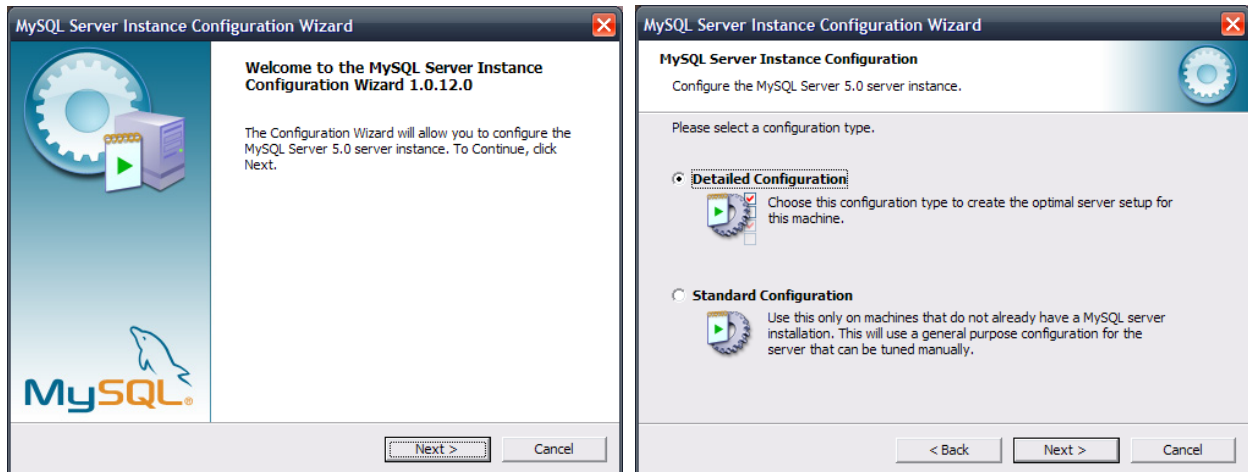

However the following links may help (Retain uses InnoDB tables):

- <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.

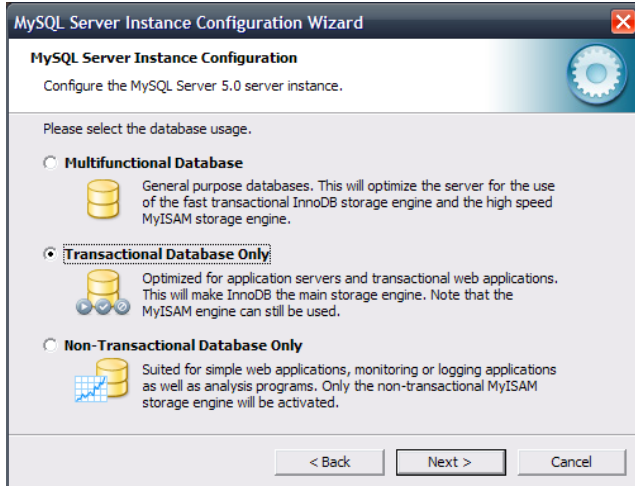
- 1) Run the installer package.
- 2) Accept all defaults.
- 3) Go to your program list and select the MySQL Instance Configuration Wizard.



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

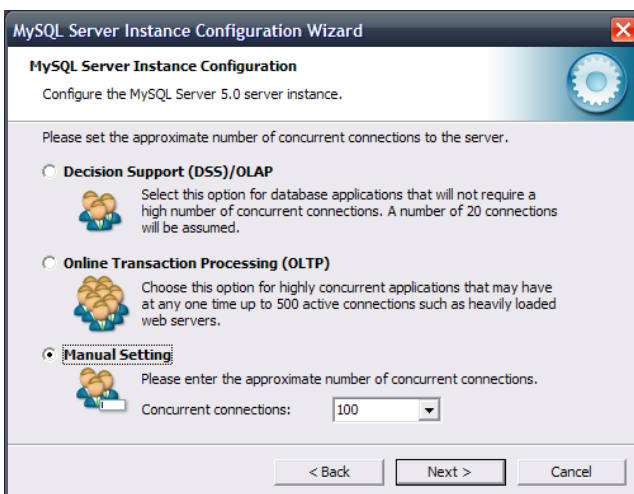



Select the appropriate configuration for the machine that is running the MySQL instance. It is recommended that you select either the Server Machine or Dedicated MySQL Server Machine setting to provide enough memory for MySQL. This will perform basic memory and CPU optimizations for your system.




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 Administration Guide.)

Choose utf-8 for the character set.



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



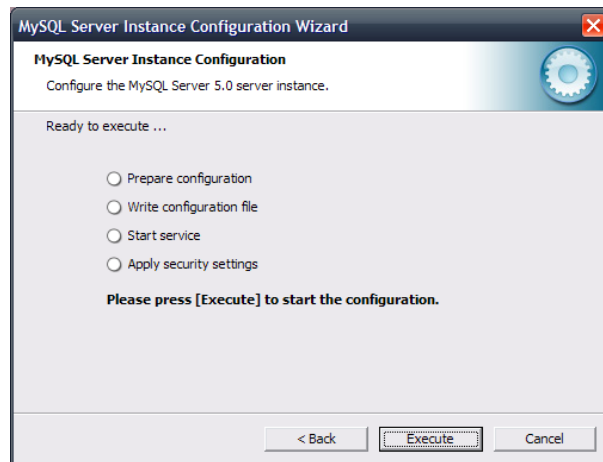
Enter the root passwords as prompted and configure remote root access as required.

WRITE THIS USERNAME AND PASSWORD DOWN!
The administrator username and password must not be lost to maintain functionality and access to the MySQL system.

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_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, 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
- Retain database name

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

SQL Server 2008 R2

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)
5. Ensure TCP/IP is enabled - Run SQL Server Configuration Manager
6. Expand SQL Server Network Configuration
7. Protocols for MSSQLSERVER
8. Enable TCP/IP
9. Run SQL Server Management Studio, connect to Server
10. Under Security/Logins, right click and select "create Login, (for example dbuser) that uses SQL Server authentication. Assign a password.
11. Right click on Databases, and create a new database named retain and assign dbuser as the owner.
(There are other ways to grant the permissions but this is easiest)

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 dbuser) that uses SQL Server authentication. Assign a password.
6. Right click on Databases, and create a new database named retain and assign dbuser as the owner.
(There are other ways to grant the permissions but this is easiest)

ORACLE 10/11g

We support the usage of Oracle 10 and 11 on all platforms. However, we do not support the installation, maintenance, or tuning of Oracle. **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\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";
pool 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.

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:

a. 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='*'
```

b. 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
```

- c. Restart PostgreSQL to activate these changes.

NOTE:

The most recent PostgreSQL offers a downloadable enterprise tuning wizard, which is useful to get tuning suggestions for your installation.

The PGAdmin windows tool may be easier than the Command line, adapt instructions as needed.

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 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 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.

Java Options

- Set the Max PermGen Size to 256M
Add 'XX:MaxPermSize=256m' to the JAVA OPTIONS screen

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

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.

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 Tomcat, launch tomcat.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.
(Tomcat is pre-tuned to 256MB initial and 1024MB max memory with a 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.

Appendix C: 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 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 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.

Appendix D: Uninstalling Retain

Should you wish to uninstall Retain, use the standard windows uninstallers.

Appendix E: 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"
```

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.

5. Create a directory c:\ssl, and cd to Apache2.2/bin directory in a DOS window.

6. Generate a server key

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

(Do NOT ever lose the password)!

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

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

8. 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

Appendix F: Windows Server 2008 and 2008 R2; IIS 7, 7.5

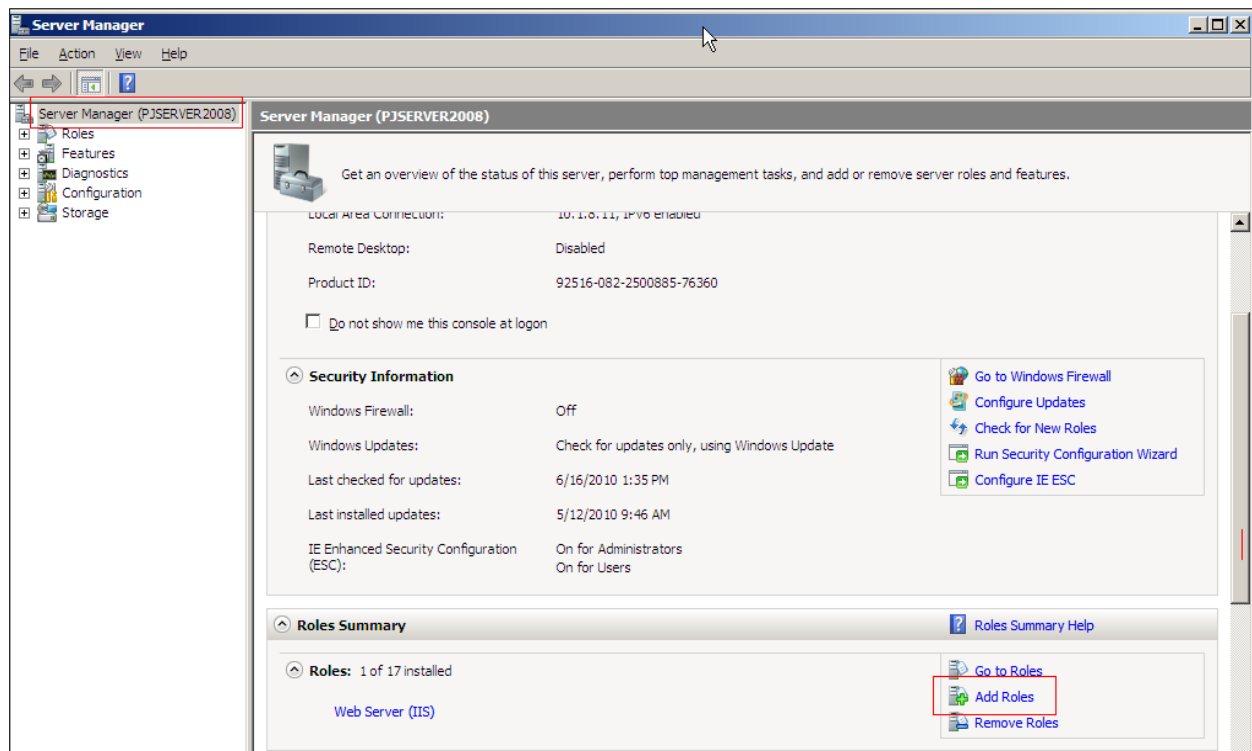
There are specific services that are required in order to integrate with Retain. Also, the Retain installation must be done in a specific order for embedded Tomcat and IIS to connect to one another. Follow these steps to successfully install Retain using IIS:

1) Install SQL Server and setup a retain database. (For further instructions on setting up the Retain database please consult the Retain documentation).

2) Install JAVA 1.6 JDK. Use the link provided to download the JAVA component.

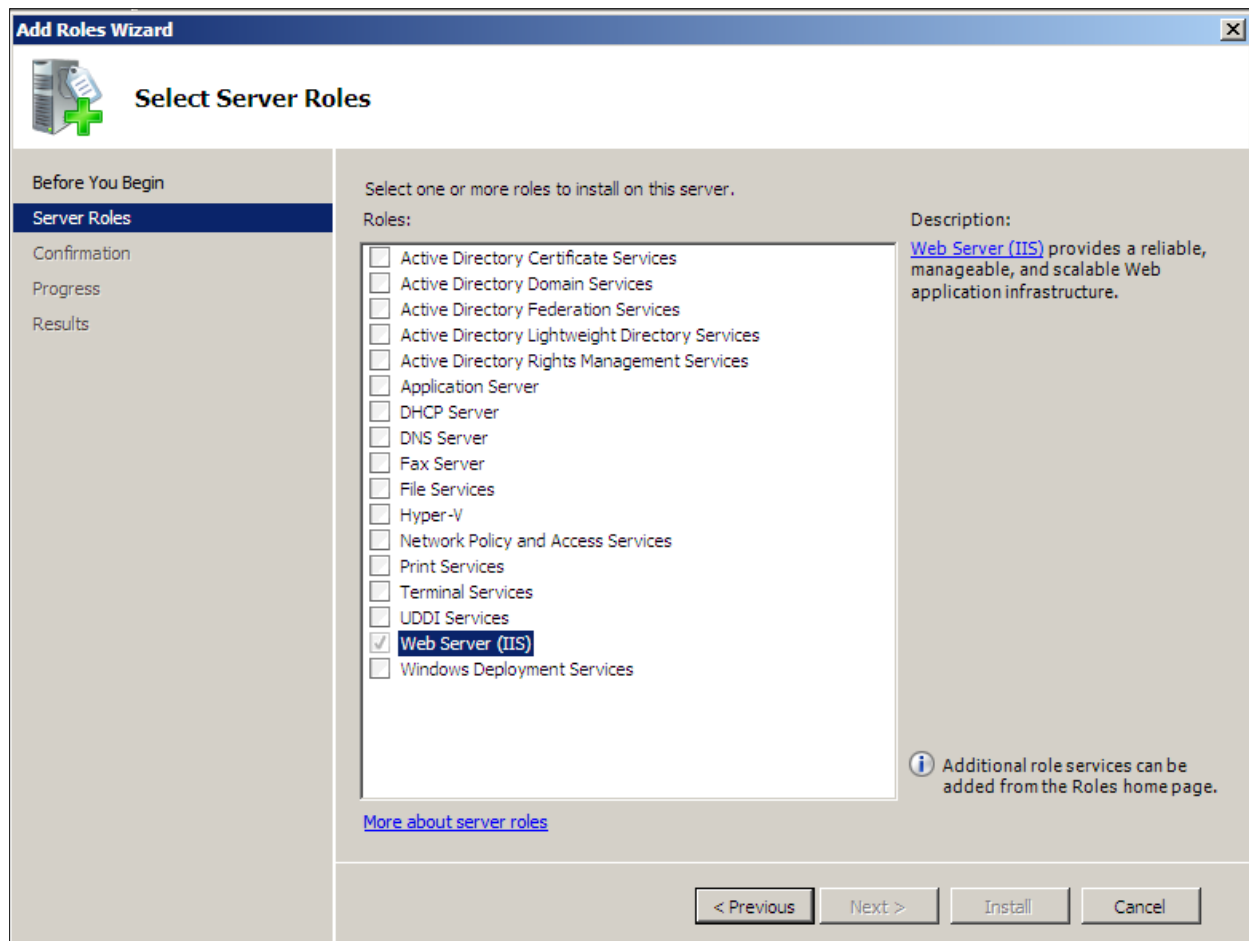
<http://download.gwava.com/jvm/jdk-6u26-windows-x64.exe>

3) Install IIS. Click **Start -> All Programs -> Administrative Tools -> Server Manager**. In the **Server Manager** window, scroll down to **Roles Summary**, and then click **Add Roles**

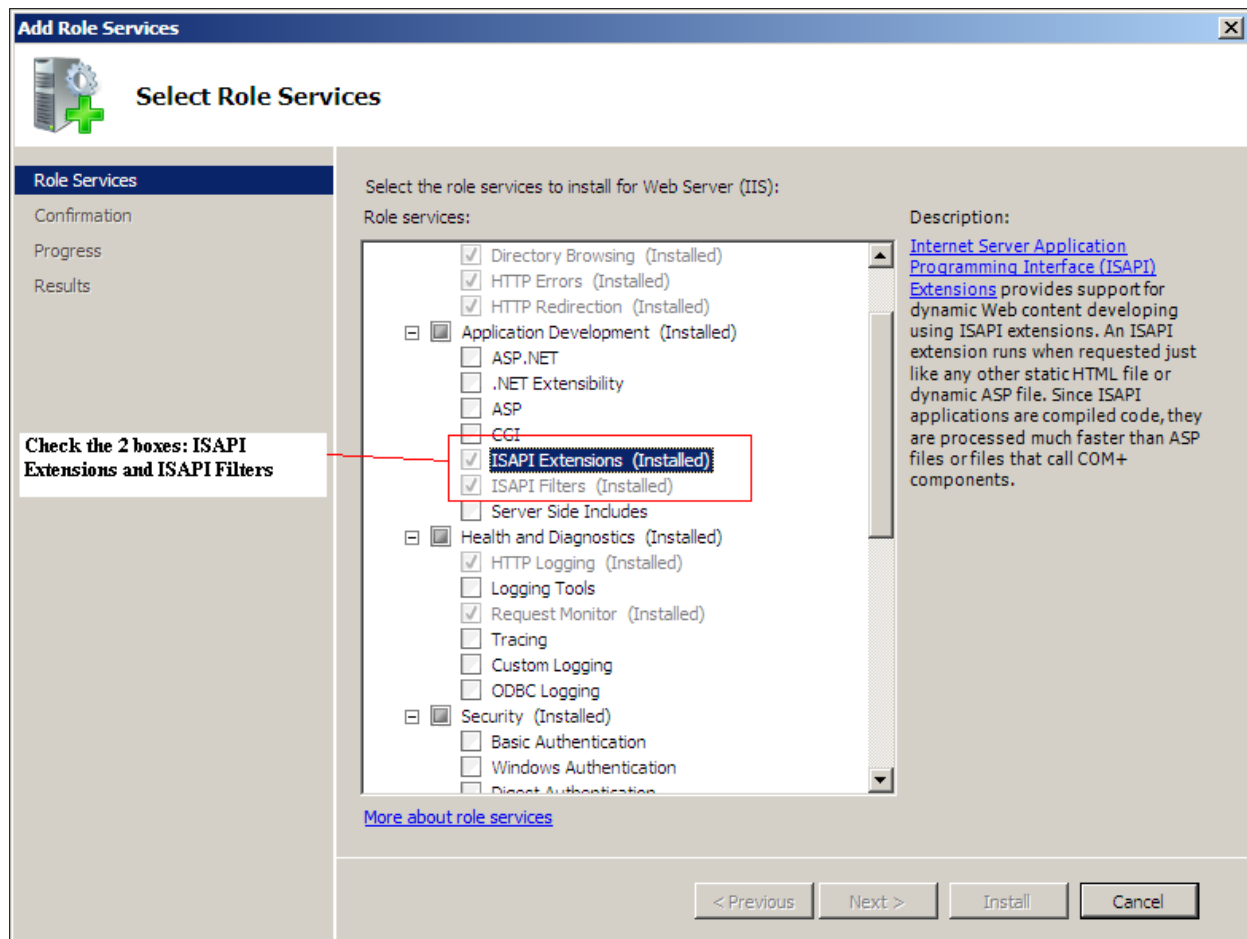


4) Click **Next** on the following page.

5) Check the box next to Select Web Server (IIS). Then click **Next**



6) In Select Role Services check **ISAPI Extensions and ISAPI Filters** found under the **Application Development** category. Leave everything else as default. Click **Next**



7) Click **Install** to install the Web Server (IIS)

8) 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. Point to your Java 1.6 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\jdk1.6.0_26

9) Install Tomcat by clicking on "Setup Tomcat for Windows.exe" and follow the prompts.

10) Run RetainInstall.exe