GWAVA Inc.

GWAVA® Retain™

Linux Install Guide

For version 2.5

Contents

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

Tomcat 5.5.27 Embedded With Retain

Retain ships with embedded Tomcat 5.5.27. Retain 2.5 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:	/opt/beginfinite/retain/tomcat
Log Directory:	/opt/beginfinite/retain/tomcat/logs
Start Script:	/etc/init.d/retain-tomcat5 start | stop | restart | status
Configuration file:	/etc/opt/beginfinite/retain/tomcat/j2ee
Memory Tuning:	Defaults to initial memory of 256MB and max of 1024MB. (Memory should still be tuned.)

The Retain install script, RetainInstall.sh, detects whether the embedded Retain-Tomcat installation is present or not, and then prompts you to install the embedded Retain-Tomcat if it is missing.

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

The hardest part about installing Retain is preparing the server(s).
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 is supported, but requirements are still in full effect and must be observed.

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

Minimum System Requirements:

- Dual-Core 3 GHz or better Processor
- 160 MB Ram free for Retain Worker
- Apache 2.2
- Supported SQL Database (Listed below)
- Sun Java 1.5 JDK 64-bit (32-bit platform supported for the Worker only.)
- As much storage space as mail 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, 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, 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
- Workers installed remotely on each post office in mail 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 Linux Versions:

- Open Enterprise Server 2 Linux (64-bit)
- SUSE Linux Enterprise Server 10 SP2 (64-bit)
- SUSE Linux Enterprise Server 11 64-bit
- 32-bit versions are only supported for the worker. (RetainServer is not supported on 32-bit.)

Microsoft Exchange Requirements:

- Exchange server 2007 SP 3
- Exchange server 2010 SP 1 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 Server2005, 2008, 2008 R2
- PostgreSQL 8.2+ , 8.3, 8.4

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. Retain dredges messages from the mail 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 server to collect data. For Exchange, Retain connects over HTTPS using Exchange Web Service Impersonation or Delegation to collect mail from all mailboxes. For GroupWise, Retain connects to the Post Office via SOAP and a trusted application key to access mailboxes and collect data. 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.

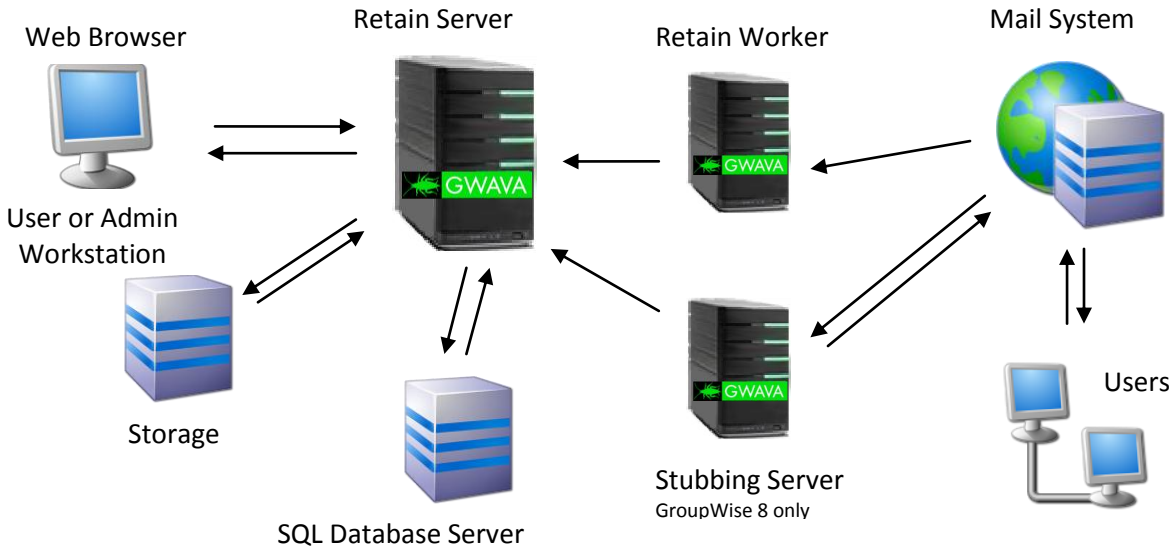

- ***Retain Server***: This is the heart and soul of Retain. All functions are controlled from the Retain Server. The server also manages the Retain Worker and stores data in the database server. There is only one Retain Server per system.

- ***Retain Worker***: There is at least one per system. The 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. Workers may exist on the same physical machine as a 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.

- Microsoft Exchange and/or Novell GroupWise
- Apache Web Server
- SQL Database for storage
- Java 1.5 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.

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 mail 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 mail 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 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 Worker can be placed on the same box as the Server for communication performance reasons. (More than one Worker sharing the same physical box as the Server is unsupported.)

Mail System Address Book

Retain gathers mail from known users. Users known to the mail 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 mail 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.

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 General Admin guide)
- (optional) Retention and/or Smart Purge flags enabled.
- It is recommended to run retention jobs during off hours to avoid overloading the mail 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 mail system.
- Address Book caching is recommended every 24 hours. (default)
- Setup is covered in the General Guide appendix.

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. See tuning Tomcat section, below.

Retain Installation

SuSE Linux Enterprise Server 10, 11

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

- Apache 2
- SQL server with appropriate database created (Unless MySQL will be utilized, see notes below)
- For installation targets without an internet connection, the latest Sun Java JDK 1.5 must be installed.

SLES 10 and 11 feature Apache 2.2x. In YaST, search for apache2; make sure Apache 2.2x is installed. Go to Network services, and activate HTTP SERVER to run on boot. ***Java 1.42 is NOT SUPPORTED.*** *The Retain install will automatically download and install the appropriate Sun JDK 1.5 if the server has a connection to the internet. If the server does not have an internet connection, you must install Sun JDK 1.5 manually.*

****** NOTE: BE SURE YOUR SERVER'S CLOCK IS ACCURATE BEFORE YOU BEGIN! USING NTP IS RECOMMENDED TO KEEP YOUR CLOCK ACCURATE.***

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

The Retain Install will complete the following tasks:

- If running/installing MySQL on the Retain Server, MySQL will be configured and a database created.
- SUN Java JDK 1.5.0.22 will be checked/installed. (32-bit (worker only) or 64-bit as appropriate)
- Apache and Retain-Tomcat will be stopped
- Retain will be unzipped into /opt/beginfinite/retain.
- The data will be stored in /var/opt/beginfinite/retain by default.
- Backups of modified files go in /opt/beginfinite/retain/backup
- The Apache integration file is in /opt/beginfinite/retain/ws
- Apache and Tomcat are made to restart and their runlevels are altered to run automatically on system start.

Install Procedure on Linux

(IF MySQL is the SQL server of choice, a installation and configuration wizard is provided which installs, configures, and creates the appropriate database for Retain. All that is required is to point to a downloaded install .rpm for the MySQL server and client packages, then follow the prompts. For more information on MySQL installation, or configuration and creation of an appropriate database for Retain, see SQL Servers in the Appendix.)

1. On the SQL database of choice, create a database for Retain and assign a user.
2. Unzip the install files to a work directory. (Your distribution directory, e.g. /root/retain)
3. Open a terminal session, such as Gnome terminal or Xterm.
4. Navigate to the distribution directory containing the unzipped install files. (e.g. cd /root/retain)
5. Install Retain-Tomcat and Retain. Enter the following commands:

`chmod +x *.sh+` (This command flags the install scripts to be executable)

`./ tomcatinstall.sh` (This command installs Retain-Tomcat and removes the Retain hooks from the standard Tomcat installation. It is recommended to disable the current Tomcat installation; you will be instructed to do so at the conclusion of the tomcatinstall.sh script. Consider removing the old Tomcat installation via YaST.)

--(optional)-- If you are not removing a previous, existing installation of Tomcat, the following command removes the previous instance from running and allocating system resources.

`chkconfig --del tomcat` (chkconfig –del tomcat5 on SLES 10, 11) (Also see the memory tuning section to return the general Tomcat tuning back to normal.)

`./RetainInstall.sh` (This command initializes the Retain installation. Move on to step 6.)

6. Read and accept the license agreement.
7. --Java is checked-- If the Java version is not Sun Java JDK 1.5.0.22, the appropriate version will be downloaded and installed. (If the system has run a 32-bit Java 1.5 version on a 64-bit server, the installer will not detect a mismatch. We strongly recommend the 32-bit version is uninstalled and the 64-bit SUN Java is installed in its place.)
8. Retain component roles. Determine whether Retain Server, Retain worker, or both will be installed.
9. Determine whether to install the Stubbing Server. **To use stubbing with a GroupWise system, it is required that the GroupWise system be at least 8.0 SP1. Previous versions of GroupWise do not contain critical stability fixes.**
   --If you have a current Retain System installed and running, you will be prompted to upgrade or overwrite your current system.--
10. Specify the SQL database server. If MySQL, and the database will be running on the same server as the Retain Server installation, the installation wizard will run which will, with prompts, install, (dependent on supplied install files), and configure MySQL, (with Administrator name and password, database storage location, database name, and access user account and privileges).
11. Installation completes and returns the prompt. Check the status of Apache 2 and Retain Tomcat. (They should both be 'running'.)

```
rcapache2 status
rcretain-tomcat status
```

You are now ready to continue on to the configuration tasks. See the General Admin guide.

Custom Storage Path Rights

During the configuration, you choose a storage path for the archived mail. Retain requires read and write access and permissions to the storage path. If you are specifying a storage path other than the default (/var/opt/beginfinite/retain) then you need to grant rights and ownership to Retain.

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

l /opt/beginfinite/retain

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

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

Create the directory structure for the desired storage location, and then grant the permissions and ownership to the Retain user.
(mkdir <desired directory strucuture>)
(chown –R <user> <storage directory>)
(chmod –R 744 <storage directory>)

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

mkdir /retain/archives
chown –R tomcat /retain/archives
chmod –R 744 /retain/archives

```
vm-sles10-retain:/retain # mkdir /retain/archives
vm-sles10-retain:/retain # chown -R tomcat /retain/archives
vm-sles10-retain:/retain # chmod -R 744 /retain/archives
```

Now check the directories for ownership.

l /retain

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

Tomcat now has rights to the storage location.

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 – ***WARNING: MySQL 4.x is provided with this version and it is installable via YaST. Do NOT use it; it is NOT compatible with Retain and will corrupt your data if used. Uninstall MySQL 4, and get MySQL 5.5x*** MySQL 5.5 needs to be downloaded and installed.

MySQL Quick Start Guide

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

To use MySQL with Retain, you will need to do the following:
1) Download MySQL
2) Install it to the server of your choice
3) Log in and create a database for Retain to use
4) Create a user with rights for Retain
5) Retain will do all the rest

Getting MySQL

Although SLES 10, and 11 come with MySQL 5.x as an optional package to be installed by YAST, the version supplied by Novell is out of date and prone to data corruption. It is STRONGLY recommended you download and install the latest RPMs from MySQL. For MySQL to be useful, three packages need to be installed: the Server, Client, and shared client libraries.

- Go to http://www.mysql.com
- Get MySQL 5.5

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

NOTE: 5.1 and 5.5 are the currently tested and supported versions of MySQL, other versions may work, but are not supported by Retain.

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/

MySQL Installation Via the Installer Script

Retain 2.5 comes with automated MySQL installation and database creation scripts, mysqlinstall.sh and mysqlsetup.sh. To use the installation scripts, both the MySQL server and client packages must be downloaded to a known location. The scripts are run during Retain Install *if* the MySQL database server is selected *and* the MySQL server will be housed on the same server as Retain. If the install initiates the scrips, they will be run back to back and do not need to be initiated manually.

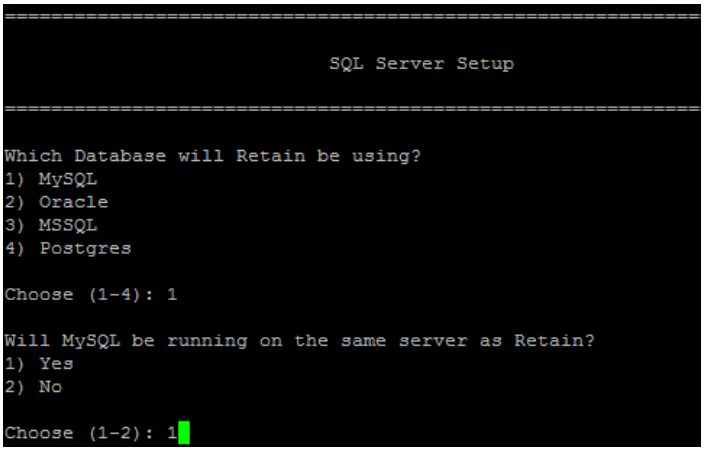

If the MySQL database is to be housed on the same server as Retain, it is recommended to run the install and configuration through the install script, but it is not required or necessary. The Database question will not change database choices in the Retain Server Configuration and setup detailed in the administration manual.

MySQL Install Script

If the MySQL database system will be run on a separate server than Retain, the install and database setup may be run manually by copying both MySQL scripts and running mysqlinstall.sh on the desired server.

The installation script will require the location of the MySQL source files of both the server and client packages. The script will search the provided location for the installation files, detect them, and automatically install the packages.

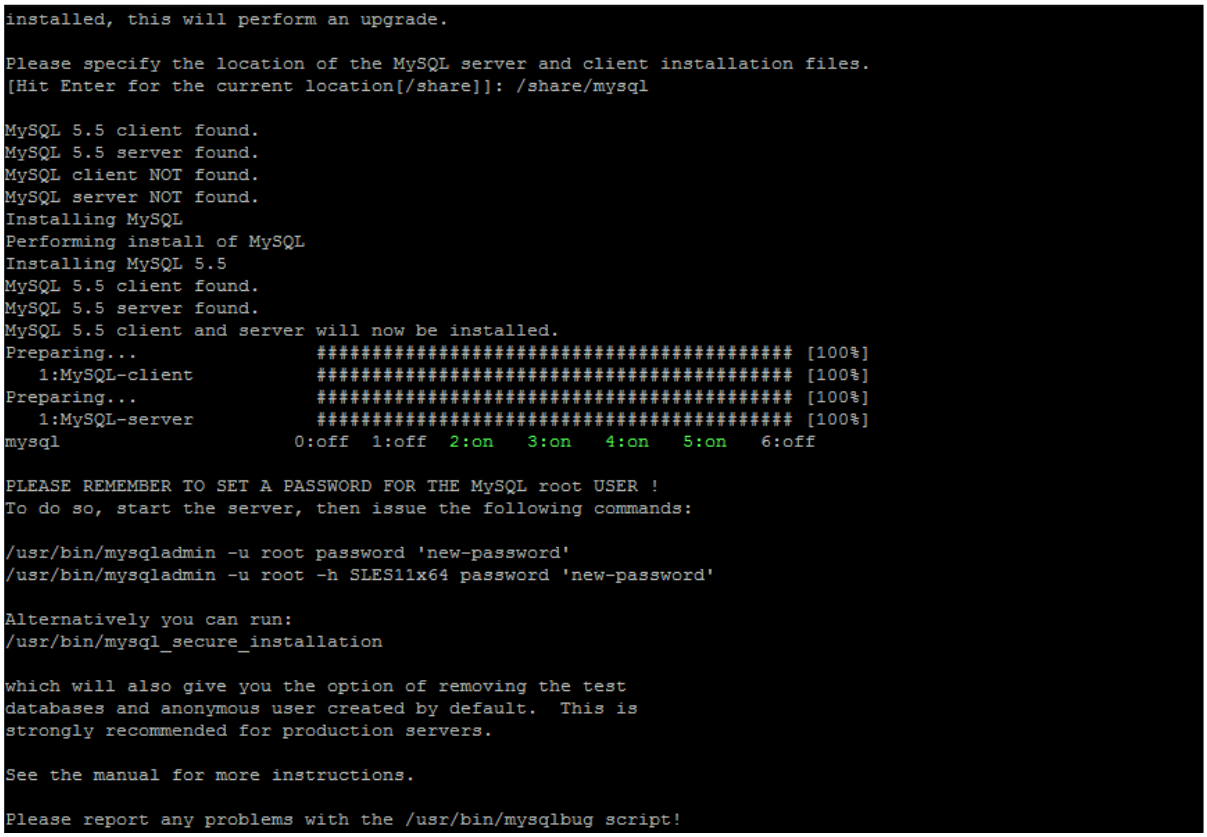

After completing, if run manually, the scrpt returns to the prompt. The script ends with instructions to manually configure the MySQL server with a root password and user. If manual configuration will be used from this point on, the commands and advice should be followed, however, if the mysqlsetup.sh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

MySQL Setup Script

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

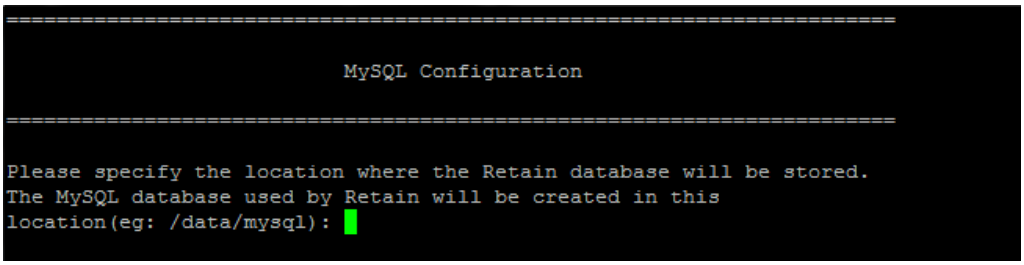

The destination folder must be specified, as there are no defaults.

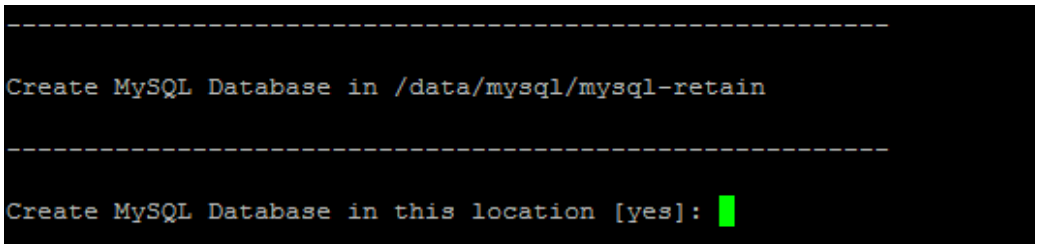

Confirm the location specified and displayed by typing 'yes' and hitting 'enter', or input 'no' to change it.

Next, the script prompts for the desired root username and password. The password must be confirmed.

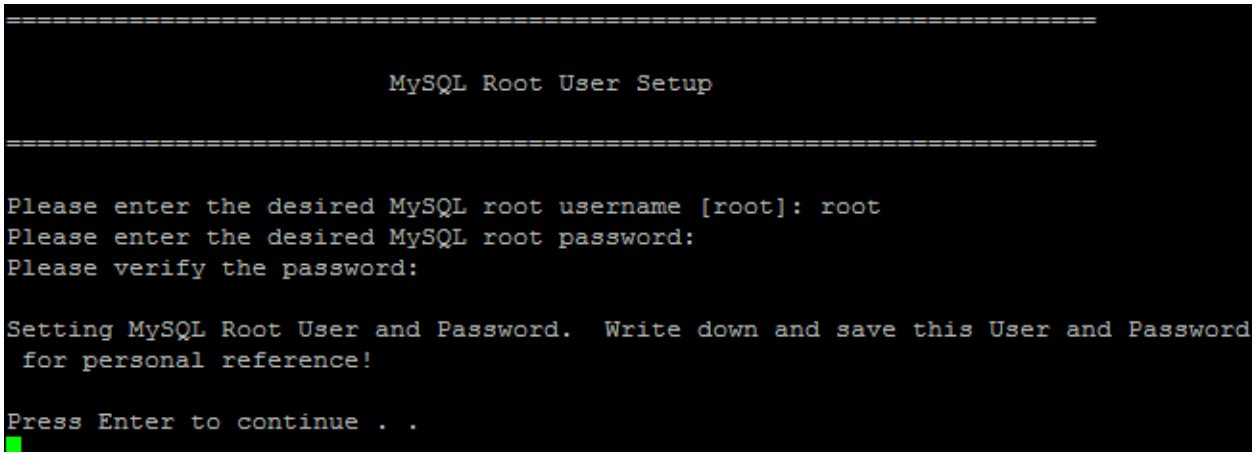

Enter and confirm the desired root username and password. This user is not the username which will be used by Retain to connect to the MySQL server, that user is setup later. This user will be the root user for MySQL.

Hit 'enter' after recording the provided password. Do not lose this password.

The MySQL database name, database user, and password specified here are all Retain specific. This database will house the Retain configuration as well as the stored messages and attachments. The database user specified will only have rights to the Retain database, and is requried by Retain to connect to and administer the database.

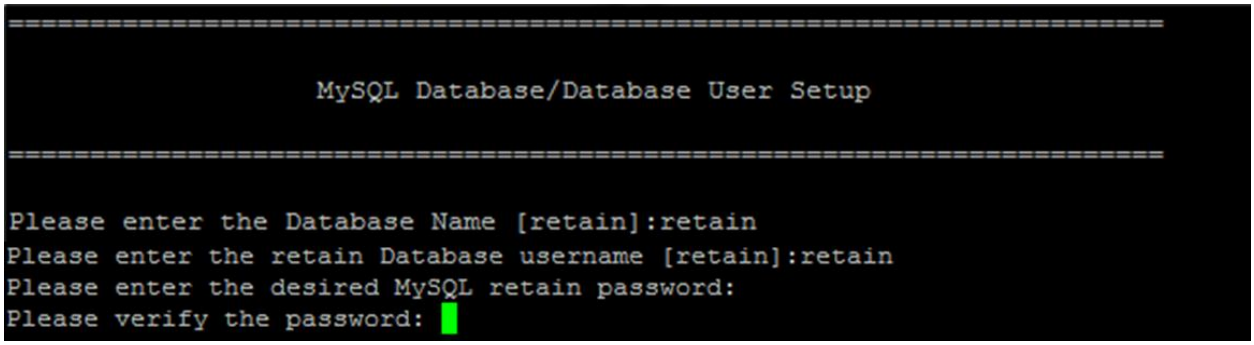

it is recommended to call the database 'retain', though the username and password may be anything desired.

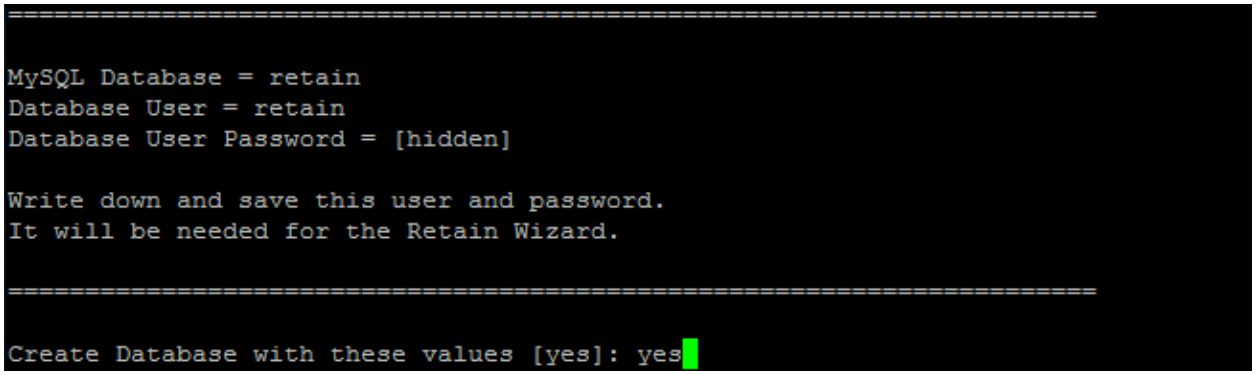

The Retain database and database user specified and listed for confirmation. Do not lose this information, as it is required for the initial Retain server setup. When the information required is acceptable, type 'yes' and hit 'enter' to allow the setup to complete. The mysqlsetup.sh script will finish,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

MySQL 5.1x Manual Installation Procedure on Linux

Applies to: SLES 10/SLES 11

1) Log in as root
2) Get MySQL 5.1x
3) Open a command-line terminal (like a DOS box in Windows)
4) Go to the directory where the RPM is.
5) For each rpm you downloaded, type
    rpm -ihv <rpmfilename>.rpm
6) Create a directory where you want your data stored.
7) Type
    mysql_install_db --datadir=<datadirectory>
8) Start MySQL:
    MySQL is a standard Linux service, with a script stored in /etc/init.d. It is configured to run automatically on restart.

The Retain 2 installer on Linux has automated the remaining steps of MySQL configuration. The following steps remain here as a courtesy for those who wish to manually create and add privileges to a database. No further configuration is needed if you are using the automated database creation provided in the installer.

9) From the command line, type
    /etc/init.d/mysql <command> --datadir=<datadirectory> where available commands are stop, start, and restart, as appropriate.

10) Create a password type:

    mysqladmin -u root password '<yourrootpassword>'

11) Log in to MySQL by typing:

mysql --user=root --password=<yourrootpassword>

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

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

Create a new database for Retain

MySQL

Applies to: All platforms

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

2) Enter the command:

```
CREATE DATABASE retain DEFAULT CHARACTER SET 'utf8' DEFAULT COLLATE 'utf8_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, enter this command:

```
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, enter the following commands:

```
GRANT ALL PRIVILEGES ON retain.* TO 'mjb'@'%' IDENTIFIED BY 'turnip';
GRANT ALL PRIVILEGES ON retain.* TO 'mjb'@'localhost' IDENTIFIED BY 'turnip';
```

(Both commands 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.

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

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

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

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

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 retainuser) that uses SQL Server authentication. Assign a password.
11. 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 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)

ORACLE 10/11g

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. **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";
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 8.2-8.4

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.

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.

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.

For the embedded Tomcat (5.5.27), which ships with Retain, Edit the following configuration file:
/etc/opt/beginfinite/retain/tomcat/j2ee

Retain-Tomcat is pre-tuned to use 256 MB initial memory, with a max of 1024 MB. This will not be enough for the server.
Edit the following line with the appropriate parameters for your system:

CATALINA_OPTS="-Xms256m -Xmx1024m -Xss160k"

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

By default, the embedded 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.

UPGRADED SYSTEMS

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

To return memory tuning to default revert the edits you created or restore a backup configuration file, (if you created one). If not, default memory settings are explained below.

SLES 10,11:

1) Edit **/etc/sysconfig/j2ee**
2) Find and change the CATALINA_OPTS line back to the original setting: (Your line may look different depending on the memory values you set.).
3) **CATALINA_OPTS="-Xms512m -Xmx1024m -Xss160k"**

Change it back to the original setting:

CATALINA_OPTS=""

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 mail systems. Different mail 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 mail 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 mail 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 mail system, it is nearly impossible to establish a baseline for Retain 2010 storage needs. Retain may tremendously decrease the needed size to archive an Exchange 2010 system, or, depending on system size and implementation, it may not significantly decrease the needs of the current system. Though Retain will require additional space to continue archiving mail, the first initial archive job will not exceed the size of the current mail 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 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. Can be: on Linux /opt/beginfinite/retain, on other platforms, stored under storage path

5. Remove the RetainWorker.xml and RetainServer.xml from Tomcat.
   a. Linux Tomcat5 /usr/share/tomcat5/conf/Catalina/localhost

6. Remove the directories under tomcat/work (but NOT work itself)
   a. Linux: `rm –r /usr/share/tomcat5/work`

7. Remove references to Retain 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: Single Sign On WebAccess Remote Install, or Manual Install for NetWare

Netware does not satisfy the requirements to successfully run the WebAccess Single Sign On module for Retain. To install manually, see the instructions found at the following link:
http://support2.gwava.com/kb/?View=entry&EntryID=1608