Maple T.A. System Administrator's Guide

Maple T.A. System Administrator's Guide

Copyright

Contents

1 Introduction

1.1 Introduction

This guide has been written for IT professionals who install and support Maple T.A. systems in colleges and universities. It addresses installation, configuration and trouble-shooting. Its goal is to provide the information necessary to perform these tasks effectively and efficiently. This guide is not a "User's Guide." In particular, it does not address the use of Maple T.A. for teaching and assessment. For that, refer to Maple T.A.'s help system.

1.2 Terminology

Throughout this document, we refer to several file system folders using a shorthand notation.

Notation	Folder
<PostgreSQL>	The PostgreSQL installation folder. On Windows®, this is by default something like "C:\Program Files\PostgreSQL\9.3". On Linux®, the default location depends on the package used. For example, on Red Hat® Enterprise Linux 6 with the default yum install it is /var/lib/pgsql.
<Tomcat>	The Tomcat™ installation folder. On Windows, this is by default something like "C:\Program Files\Apache Software Foundation\Tomcat 7.0". On Linux, the default location depends on the package used. For example, on Red Hat Enterprise Linux 6 it is /usr/share/tomcat7 (for a Tomcat 7 installation).
<Maple T.A.>	The Maple T.A. installation folder. On Windows, this is by default something like "C:\Program Files\Maple T.A. 10". On Linux, the default is to use the folder "Maple T.A. 10" with the same parent folder as <Tomcat> above.

2 Installation

2.1 Platform Prerequisite Hardware and Software

Summary

Maple T.A. is a Java™-based Tomcat web application that uses a PostgreSQL database. Maple T.A. can run on Windows, Linux, and Macintosh®.

Server Platforms and Requirements

You should run Maple T.A. on a dedicated stand-alone server. In particular, note the following restrictions:

1. No other web applications should be sharing Maple T.A.'s Tomcat container (except for Tomcat's administrative applications).
2. No web server other than Tomcat should be running on the machine.
3. The machine should not be providing any non-Maple T.A. services to other machines. It should not be a database server, a file server, a DNS server, etc.

The following are the operating systems supported, minimum requirements, and recommended requirements for each platform:

Platform	OS Versions	Minimum Requirements	Recommended Requirements *
Windows (32-bit)	Windows XP Windows 7 Windows 8.1	• Dual core CPU • 2 GB RAM • 20 GB hard drive space	• 8-core CPU • 32 GB RAM • 40 GB hard drive space
Windows (64-bit)	Windows XP Windows Server 2008 R2 Windows Server 2012 Windows 7 Windows 8.1	• Dual core CPU • 2 GB RAM • 20 GB hard drive space	• 8-core CPU • 32 GB RAM • 40 GB hard drive space
Linux (32-bit and 64-bit)	SUSE® Linux Enterprise Desktop 11 Red Hat Enterprise Linux 6 CentOS® 6.5 Ubuntu® 12.04 LTS, 13.10	• Dual core CPU • 2 GB RAM • 20 GB hard drive space • X window system or xvfb configured	• 8-core CPU • 32 GB RAM • 40 GB hard drive space • X window system or xvfb configured
Macintosh (64-bit)	Mac OS® X 10.8, 10.9	• Dual core CPU • 2 GB RAM • 20 GB hard drive space	• 8-core CPU • 32 GB RAM • 40 GB hard drive space

* Recommended for concurrent usage of up to 600 tests peak, or general enrollment of up to 3000 students.

There is no particular reason that Maple T.A. will not run on other Linux distributions and on other versions of the previously mentioned distributions, but your mileage may vary. The preceding distributions and versions are those that Maplesoft puts through its full QA process. Maplesoft Technical Support will assist on a best efforts basis with problems encountered on other Linux distributions and on other versions of the preceding distributions.

Client Platforms and Requirements

All clients must have access to the Internet. A broadband (Cable modem or DSL) Internet connection is recommended for optimal performance.

JavaScript™ and cookies must be enabled.

The following are the operating systems supported, browsers supported, and hardware requirements for each client platform:

OS Version	Browsers	Hardware Requirements
Windows XP	Firefox®, Google Chrome™	1 GHz processor or better, 2 GB RAM or better
Windows Vista®	Firefox®, Google Chrome™, Internet Explorer®	1 GHz processor or better, 2 GB RAM or better
Windows 7	Firefox®, Google Chrome™, Internet Explorer®	1 GHz processor or better, 2 GB RAM or better
Windows 8	Firefox®, Google Chrome™, Internet Explorer®	1 GHz processor or better, 2 GB RAM or better
Windows 8.1	Firefox®, Google Chrome™, Internet Explorer®	1 GHz processor or better, 2 GB RAM or better

OS Version	Browsers	Hardware Requirements
Mac OS X 10.8 or 10.9	Firefox®, Safari®, Google Chrome™	1 GHz processor or better, 2 GB RAM or better

OS Version	Browsers	Hardware Requirements
SUSE Linux Enterprise Desktop 11	Firefox®, Google Chrome™	1 GHz processor or better, 2 GB RAM or better
Red Hat Enterprise Linux 6	Firefox®, Google Chrome™	1 GHz processor or better, 2 GB RAM or better
Ubuntu 12.04 LTS, 13.10	Firefox®, Google Chrome™	1 GHz processor or better, 2 GB RAM or better
CentOS 6.5	Firefox®, Google Chrome™	1 GHz processor or better, 2 GB RAM or better

OS Version	Browsers
iOS® 7 on iPad®	Safari®
Windows® 8/8.1 RT	Internet Explorer® (Desktop mode only)
Android™ 4.4	Google Chrome™

For Students: Instructors may require that students take a test in Maple T.A. Proctored Browser mode. To use the Proctored Browser, you must use Google Chrome or Firefox.

Installing Java on the Server

Maple T.A. requires Java 6 or Java 7. Java 6 is required for server clustering.

A Java Development Kit (JDK™) or Java runtime environment (JRE) are each suitable. The smaller JRE is adequate.

For Windows or Linux

1. Download and install the latest Java Runtime Environment (JRE) from Oracle's Java SE site: **http://www.oracle.com/technetwork/java/javase/downloads/index.html**.
   For Linux, you will typically have several choices, including installing from binaries downloaded from **http://www.oracle.com/technetwork/java/javase/downloads/index.html** and including packages from various sources, depending on your Linux distribution. We recommend using a package that installs a recent version of Java 6 as a JRE.

2. After installation of Java, set your system-wide **JDK_HOME** or **JRE_HOME** environment variable, depending on whether you installed the JDK or JRE. These variables are set to the base of the Java installation folder (the parent of the Java "bin" folder).

3. Under Windows, add Java's `bin` directory to your System path. Under Linux, this should only be necessary with a binary install, as package installs typically install the Java to be executable in such a way that it can be found through the system path.

For Macintosh

Macintosh Note: Mac OS X 10.8 and Mac OS X 10.9 do not provide a JRE by default. Follow these instructions to install a JRE on your Mac

1. To check if you have Java installed, go to **Applications > Utilities > Java Preferences**.

2. If Java is not installed, you will be prompted to install it.

3. Ensure Java is enabled. In **Applications > Utilities > Java Preferences**, ensure that a JRE is selected.

Installing PostgreSQL on the Server

Maple T.A. requires PostgreSQL 9.3.x. **Note :** After installing PostgreSQL, you must make important changes to the PostgreSQL configuration before installing Maple T.A. See *Required PostgreSQL Configuration Changes (page 4)*.

You can install PostgreSQL in one of two ways:

- By using the provided PostgreSQL installers on the Maple T.A. DVD. Navigate to the `postgres` folder of the Maple T.A. DVD and select the appropriate file for your platform.

- By downloading the latest PostgreSQL version from the PostgreSQL web site: **http://www.postgresql.org/download**:

1. Under the heading **Binary packages**, there are links for **Linux**, **Mac OS X**, and **Windows**. Select the appropriate link for your platform.

2. Click the **Download** link.

3. On the downloads page, under the latest version, click the link for your platform.

To install PostgreSQL:

1. Start the PostgreSQL installer and follow the on-screen instructions.

2. You will be asked to provide a PostgreSQL administrative username and password. Record it; you will need to use it during the Maple T.A. installation.

Note for Linux Users

1. You will typically have several choices, including installing from binaries downloaded from **http://www.postgresql.org/** and including packages from various sources, depending on your Linux distribution. We recommend installing a package with a recent version of PostgreSQL.

2. Install pgadmin (the PostgreSQL GUI database explorer tool) if you have the option, and if you plan to run a full X server (as opposed to the minimal server needed for Maple T.A., viz. xvfb).

3. A package install will often create an administrative user "postgres" without a password. You will need an administrative user with a password for the Maple T.A. install. The easiest way to fix this is to create a password for the postgres user, with psql. Something like this will work on most Linux versions:

    sudo su –c psql –U postgres

    # alter user postgres password xxxxxxxx;

    # \q

You may also be able to set the postgres password using "sudo pgadmin" if you have pgadmin installed.

Required PostgreSQL Configuration Changes

Important: Maple T.A. requires a PostgreSQL configuration change before you run the Maple T.A. installer.

To set these configurations for PostgreSQL:

1. Stop the PostgreSQL server.

2. Back up and then open **<PostgresSQL>/data/postresql.conf** for editing.

3. Find **standard_conforming_strings** settings in the file.

4. Make sure the line is not commented out (remove '#' at the beginning of the line) and set the value of **standard_conforming_strings** to **off**:

    standard_conforming_strings=off

5. Find **bytea_output** settings in the file.

6. Make sure the line is not commented out (remove '#' at the beginning of the line) and set the value of **bytea_output** to **'escape'**:

   bytea_output='escape'

7. Start the PostgreSQL server.

Verifying the PostgreSQL installation

A basic smoke test is enough. Just make sure you can log in with your superuser id (usually "postgres") with pgadmin or psql.

Required Tomcat Configuration Changes

Maple T.A. requires two Tomcat configuration changes.

1. Internationalization — If you wish to use characters not found in the basic English alphabet you need to set Tomcat's default text encoding to "UTF8" (a form of Unicode). Start the installer, and follow the on-screen instructions.

Note: Any connectors set up in the <Tomcat>/conf/server.xml file must be configured to utf-8 encoding. Specifically, the attribute URIEncoding="UTF-8" must be added to any <Connector ... /> blocks in the server.xml file that are not commented out.

2. Performance and Scaling — All but the smallest Maple T.A. installations must run Tomcat with more Java memory ("heap space") than is allocated by default.

Configure these two items as follows:

Windows

To set these configurations for Tomcat:

1. Right-click on the Tomcat System Tray icon.

2. Select **Configure**.

3. Select the **Java** tab.

4. Set **Initial memory pool** to 128 and **Maximum memory pool** to 512.

5. Under **Java Options** add the line:

-Dfile.encoding=UTF-8 .

Linux and Macintosh

To set the JAVA_OPTS environment variable for Tomcat:

• Add the following line to the **<Tomcat>/bin/catalina.sh** file

export JAVA_OPTS="$JAVA_OPTS -Duser.language=en -Dfile.encoding=UTF-8 -Xms128M -Xmx512M -XX:PermSize=64M -XX:MaxPermSize=512M"

Note: There may be many other ways to set this variable, depending on the distribution, including (but not limited to):

• /etc/default/tomcatx[1]

• /etc/init.d/tomcatx

• The .profile file for the user that runs the Tomcat daemon

• The system .profile file

[1] The "x" will depend on the version of Tomcat.

Providing an X Server under Linux

Under Linux , Maple T.A. requires an X-Server to provide a graphics environment needed to render mathematical symbols and graphs.

The minimal way to accomplish this is with a virtual frame buffer. We recommend xvfb.

A simple implementation is to insert the following into **/etc/init.d/tomcatx**, preceding the invoking of **catalina.sh**:

/usr/bin/Xvfb :1 -screen 0 1024x768x8 &

export DISPLAY=:1

If you do not start Tomcat from an **init.d** script, then devise something similar that fits your approach.

2.2 Installing Maple T.A.

Information Requested by the Installer

The installer needs the following information for software already installed:

• <Tomcat> — This is your Tomcat installation folder.

• The username and password for a PostgreSQL superuser. Usually this will be the "postgres" user.

The installer needs the following information about your environment:

• Your time zone — Be aware that Maple T.A. uses the time zone you provide at installation. The default suggested at installation is the time zone used by your server. We recommend that you use that default. It is possible to specify a different time zone than that of the server but situations where that is required are extremely rare. We recommend that a single instance of Maple T.A. run on a dedicate server. In this recommended scenario, there is no reason that the server time zone should not be set to the value needed for Maple T.A. That is, there is no reason for the Maple T.A. and server time zones to be different.

• Details as to how you want Maple T.A. to send mail: server name, from address, authentication credentials.

• Authentication through LDAP — If you want to authenticate through LDAP, configuration of the LDAP interface is done post-install. You will not need to answer any questions on it during the install.

The installer will also ask you to provide names for created items.

• The pathname for the Maple T.A. installation folder. Under Windows, we recommend the usual "Program Files" parent with the default "Maple T.A. 10" folder name. Under Linux, we recommend the parent folder of the Tomcat installation as the parent folder for Maple T.A. with the default "Maple T.A. 10" folder name.

• The name of a PostgreSQL database to be created for Maple T.A. We recommend the default name, "mapleta".

• A PostgreSQL username and password to be used by Maple T.A. to access PostgreSQL.

• A Maple T.A. username and password to be used as the principal administrator. Note that for LDAP installations, this username must be in LDAP, or you will not be able to log in with it.

Running the Installer

The installer is run in the typical way. See *Linux Considerations (page 6)* for information specific to installing on Linux.

Changes Made by the Installer

When it is run, the installer:

1. Installs a Maple T.A. installation folder in the location you choose during installation

2. Deploys a mapleta Tomcat web application in Tomcat's default webapps folder - <Tomcat>/webapps

3. Adds a user to PostgreSQL

Linux Considerations

• The installer runs as an X application by default. Specify the 'console' parameter to run it in text mode.

• The installation must be run either as root (sudo) or under the user that runs the Tomcat daemon (usually tomcat or tomcat5 or similar). This is required to ensure that the installer can assign the necessary permissions to the Maple T.A. installation files and folders.

- Ubuntu "Error listenerStart". Ubuntu's package installer may install Tomcat with so-called "Tomcat Security" enabled. This is actually Java code security[2]. Ubuntu also installs Tomcat in a "multiple instances" fashion, one aspect of which is that the default webapps directory is physically outside the Tomcat folder (usually it is installed as /var/lib/tomcat/webapps)[3]. The combination of these two things prevents Maple T.A. from functioning correctly. The symptom is that the mapleta web application does not start. Tomcat's catalina.log shows "Error listenerStart" followed by a message indicating that mapleta failed to start. The problem is caused by the fact that Maple T.A. uses Hibernate, which uses cglib, which requires elevated privileges that are not available by default in the multiple-instances T.A. installation, and which are enforced by the Tomcat Security policy. The fix is to append the following configuration setting to <Tomcat>/conf/policy.d/04webapps:

    grant codeBase "file:${catalina.base}/webapps/mapleta/-" {

    permission java.security.AllPermission;

    };

- Ubuntu java.lang.NullPointerException
  at com.maplesoft.mathdoc.font.WmiFontResolver.getFontForCharacter(Unknown Source)
  This arises from the Ubuntu default Java setting java.awt.headless=true
  Change this, in JAVA_OPTS in /etc/default/tomcatx, to java.awt.headless=false

- Various Linux distributions javax.xml.transform.TransformerFactoryConfigurationError: Provider
  org.apache.xalan.processor.TransformerFactoryImpl not found
  This can arise from incompatibilities between the versions of Java and Tomcat that are installed
  Normally this can be fixed by adding the following to JAVA_OPTS (for example, in /etc/default/tomcatx if your installation has that file).

    -Djavax.xml.transform.TransformerFactory=

    com.sun.org.apache.xalan.internal.xsltc.trax.TransformerFactoryImpl

[2] Tomcat Security is documented in the Tomcat User Guide, chapter "Security Manager". See **http://tomcat.apache.org/**.

[3] The Tomcat multiple instances model is described in the Tomcat Configuration Reference, chapter "The Host Container". See **http://tomcat.apache.org/**.

2.3 Verifying the Maple T.A. Installation

Before verifying the Maple T.A. installation, you should verify the Tomcat installation.

Maple T.A. Startup Smoke Test

1. Stop Tomcat.

2. Delete all of the log files in <Tomcat>/logs.

3. Start Tomcat.

4. Browse to your Maple T.A. home page. With installation defaults, this is http://localhost:8080/mapleta.

5. Log in to Maple T.A. with the administrator credentials that you created during installation.

6. Stop Tomcat.

7. Inspect catalina.log. The smoke test "pass" criterion is that you should see no ERROR messages and no exception stack traces.

Basic Maple T.A. Functionality Test

1. Log in to Maple T.A. with the administrator credentials that you created during installation.

2. Select the **Maple T.A. Readiness Class** by clicking on the class name.

3. You will see a page with a grid containing demonstration assignments.

4. Click **Student Readiness Test**.

5. Go through the test to ensure the server is functioning properly.

2.4 Reinstalling Maple T.A.

The Maple T.A. installer will not overwrite an existing Maple T.A. database or an existing database user (as the admin user specified during installation). Therefore, to reinstall with the same names (the normal case) you need to remove these items. This procedure requires the use of PostgreSQL tools. See *Using PostgreSQL Tools (page 39)* if you are not familiar with these tools.

The reinstallation steps are:

1. Undeploy Maple T.A. from Tomcat.

   a. Log in to the Tomcat Manager application (see "Appendix B – Using the Tomcat Manager Facility ").

   b. Click the mapleta "Undeploy" link.

   c. When Tomcat has completed the Undeploy, the mapleta entry will disappear from the Application list in the Tomcat Manager display.

   d. If the <Tomcat>/webapps/mapleta folder has not been completely removed by the Undeploy, remove it manually.

2. Uninstall Maple T.A.

   a. Move the file <Maple T.A.>/license/license.dat to a temporary location. You will use it in step 5.

   b. Under Windows, use Add/Remove Programs from the Control Panel to remove Maple T.A. Under Linux, run the uninstall binary in the <Maple T.A.>/bin folder.

   c. If the <Maple T.A.> installation folder has not been completely removed by the uninstaller, remove it manually.

3. Remove the mapleta database from PostgreSQL.

   a. Using the pgAdmin tool:

      i. Right-click on the mapleta database node in the pgAdmin tree view.

      ii. Select Delete/Drop.

   b. Alternatively, using psql:

      i. Issue the SQL command 'drop database mapleta;'.

4. Remove the mapleta user from PostgreSQL.

   a. Using the pgAdmin tool:

      i. Click on the Login Roles node in the pgAdmin tree view.

      ii. Right-click the mapleta user.

      iii. Select Delete/Drop.

   b. Alternatively, using psql:

      i. Issue the SQL command 'drop role mapleta;' (replacing "mapleta" by the name of the mapleta user, if it is different).

5. Install mapleta.

   a. Installation will now proceed in the same manner as a "fresh" install.

6. Restore license.dat

   a. Move license.dat (which you saved in step 2.a) to <Maple T.A.>/license/. This avoids the need to re-activate.

3 Post-Installation Configuration Options

3.1 The System Admin User Interface

The System Admin user interface consists of:

1) Four links situated in the top right corner; My Homepage, Help, your profile page, and Logout.

2) Three menu options: Class Manager, System User Manager, and System Admin.

The remainder of this chapter will focus on the System Admin menu options. For information on the Class Manager and the System User Manager menu options , see chapter 4 and 5, respectively, of the Maple T.A. System Administrator online help,

The System Admin menu options allow system administrators to change properties and settings of the Maple T.A. install. Additionally, from the user interface, system administrators can manage other aspects of the Maple T.A. system, for example, archiving, usage reporting and logs.

The System Status Page

The System Status page provides a summary of connections and recent usage.

To view the System Status page:

1) From the **System Admin** menu, select **System Status**.

2) The System Status page is displayed.

The System Settings Page

The System Settings page itself is divided into eight distinct sections: System Settings, Mail Settings, Authentication Settings, Maple Settings, Custom CSS Styles, Custom Head HTML, Custom Footer HTML and IP Address Groups Settings.

To access the System Settings page:

1) From the **System Admin** menu Select **System Settings** to view and configure system settings.

The System Settings Panel

The **System Settings** panel provides an interface for system administrators to change the path of the Maple T.A. installation, limit the size of uploaded files, control the number of users allowed on the system, and set the time between log file generation.

Property	Default	Comments
Gateway Home Path	None	Maple T.A. installation directory. This directory is set during installation and can be changed through the system settings page and directly in the system_properties table of the database.
School Name	None	School Name is a required field.
Update School Logo	None	School logo to display in the top-right of the header of every page, including the login page. If none is provided, a Maple T.A. logo is displayed. The recommended size of the school logo should be 140 x 60 pixels.
Time Zone	None	Time Zone can only be entered through the drop down list box.
Allowed File Extensions	htm, html, jpg, gif, png, pdf, txt, m4v, mw, mla, wmv, avi, ggb, wav, mp3, swf, lib, ind, jp	Allowed file extensions for files you wish to upload through the Class File Manager in a class.
Max Users	None	The number of student users allowed on the system concurrently.
Show Student ID	True (checked)	If checked, students will see their student ID displayed on the Assignments page in the top center of the page.
Anonymous Homepage	True (checked)	Allow anonymous users to access system homepage. For more information, see the

		section on Allowing Anonymous User Access below.
Max Upload Size (kB)	2000	Maximum size, in kilobytes, of an uploaded file. Uploaded files include question banks and course modules.
Usage Sleep Log (mins)	20	Time limit, in minutes, to wait between generation of usage logs.

Allowing Anonymous User Access

A class can be set up to allow anonymous access, so that anonymous users can take practice tests for testing or self-evaluation purposes. For a class to have the option to allow anonymous access, anonymous access must be enabled at the system level first.

To enable system level anonymous access:

1. From the **System Admin** menu, select **System Settings**.

2. In the **System Settings** panel, select the **Anonymous Homepage** check box.

3. Click **Submit.**

To enable class level anonymous access:

1. From the **Class Manager** menu, select **Class Search**.

2. Search for the class that you want to enable anonymous access to.

3. Click the class name from the **Active classes in the system** panel.

4. In the **Class Details** panel, click **Class Info**.

5. Click **Edit**.

6. Select the **Anonymous Hompage** check box.

7. Click **Submit**.

The Mail Settings Panel

The **Mail Settings** panel of the **System Settings** page provides an interface for system administrators to configure email settings for Maple T.A.
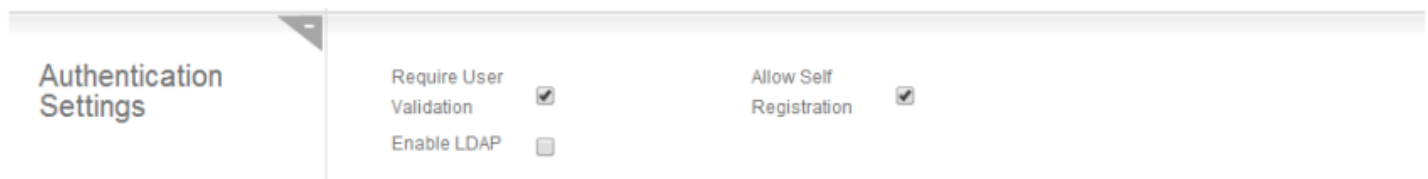

Property	Default	Comments

MX Lookup	Unselected	Find and use MX records for the given domain.
Domain Name	None	If the property "MX Lookup" is selected, you will need to provide the domain name.
Generic	Selected	Use a mail server from the server list.
Servers	Localhost	If the property "Generic" is selected, you must provide a single server or a list of servers separated by a semicolon (;).
From Address	Defined during installation of Maple T.A	This is the address that will appear in the from field of any email sent.
User Name	None	The user name used to connect to the mail server (if needed).
Password	None	The password used to connect to the mail server (if needed).
Format	Text	Determines whether email will be sent as HTML or text only.
Test Connection	None	The test connection button takes the settings that are in place and attempts to connect to the mail server. The user is informed whether or not the test was successful.

The Authentication Settings Panel

The System Administrator can control user authentication and validation through the **Authentication Settings** panel of the **System Settings** page.

Property	Default	Comments
Enable LDAP	False (unselected)	Selecting the Enable LDAP box, means that Maple T.A. LDAP authentication needs to be configured through the ldap.properties file. For more information about Maple T.A. LDAP authentication, see *Maple T.A. LDAP Authentication (page 26)*.
Require User Validation	True	Values are "true" or "false". If "true", users logging in for the first time will be required to confirm their profile details (first name, last name, email, etc).
Allow Self Registration	False	Values are "true" or "false". If "true", users will see a link on the login page that allows them to register themselves. This feature is normally used only by Maple T.A. installations that are oriented towards placement testing.

The Maple Settings Panel

The local pool of Maple servers is a collection of instances of command-line Maple that are started as needed by Maple T.A. to compute versions of questions (for example, randomized questions) and to grade questions. This pool is configurable through the System Settings page. To access the System Settings page, select **System Settings** from the **System Admin** menu.

By default, Maple T.A. uses the local pool of servers described in the paragraph above. It is also possible to use Maplesoft's MapleNet product as an alternative to the local pool. There are two reasons for doing this:

1. Provide a wider range of question features including maplets and gridlines on plots.

2. Scalability – the option of load balancing by running Maple T.A.'s Maple computations on a separate server.

Maple T.A.'s use of MapleNet is configured by properties in the System Admin user interface, under the Maple Settings section.

If all of the above properties are configured correctly, Maple T.A. will use MapleNet. Otherwise, Maple T.A. will use its local pool.

You can confirm that Maple T.A. is using MapleNet by inspecting mapleta.log. The following message will be issued at startup:

INFO This Maple T.A. instance uses MapleNet Services

Note that you will have to increase the log level for mapleta.log to INFO, from the default WARN, to see this message. See *The Default Log Level (page 22)*.

To configure MapleNet itself, refer to the MapleNet documentation. In particular, be aware that when Maple T.A. is configured to use MapleNet as the back-end math engine, and questions need access to files or custom libraries, the following may be required:

1. MapleNet may need to be run without security option "-z" to enable general access to the file system.

2. If Maple T.A. and MapleNet are running on two different servers, the absolute path to the library file from the Maple T.A. server needs to be identical to that seen by MapleNet on the other server. This can be done by copying the required files, or by using a shared network file system.

Maple Settings

Field	Value	Field	Value
Use MapleNet	☐		
Server	199.71.182.78	Port	80
User Name			
Password			
Maple Path	/usr/local/hosted_apps/Maple1801.20140522/bin/maple		
Min Servers ?	1	Max Servers ?	10
Max Restarts ?	250	CPU Limit (s) ?	120
Retry Delay (s) ?	20	Time Limit (s) ?	20
Plot Driver	GIF ▾	Size Limit (kB) ?	100
Plot Height (px)	360	Plot Width (px)	480
Enable Math App Question	☑		
Shared Math App File System	☐		
Server	199.71.182.78	Port	80

Property	Default	Comments
Retry Delay(s)	120	Number of milliseconds to wait between requests for a Maple server instance, when all servers in the pool are busy.
Max Servers	5	Maximum number of Maple servers in the pool.

Property	Default	Comments
Min Servers	1	Minimum number of Maple servers in the pool.
CPU Limit(s)	120	Maximum number of cpu seconds that can be consumed by a Maple instance in the pool, before it is aborted (assumed to be a "runaway").
Max Restarts	250	Maximum number of maple commands that will be sent to a Maple server before it is terminated and replaced with a fresh server.
Plot Width (px)	400	Plot width in pixels.
Plot Height (px)	800	Plot height in pixels.
Plot Driver	gif	Graphics format used to render plots. The choices are gif and jpeg.
Server	<None>	The domain or subdomain address of the MapleNet server, e.g. maplenet.myu.edu. MapleNet can run on the same machine as Maple T.A. In that case you would specify MAPLENET_SERVER = localhost.
Port	<None>	The port number listened to by MapleNet on its server, e.g. 8180.
User Name	<None>	The username configured for MapleNet.
Password	<None>	The password for MAPLENET_USER.

The Custom CSS Styles Panel

You can override the default styling in Maple TA by providing CSS styles. This allows you to customize the look of Maple T.A.'s fonts, colors, borders, etc.

Custom CSS Styles

Include custom CSS styles that will be applied to every page in Maple T.A. Use this to override Maple T.A.'s default styling.

The Custom <head> Panel

You can place custom code in the <head> tag of every Maple T.A. page. The custom code can be extra CSS, JavaScript files, custom scripts, etc.

Custom <head> HTML

Include custom HTML in the <head> tag of every page. Allows for inclusion of custom script tags or other includes.

The Custom footer HTML Panel

Code added to the Custom footer HTML panel is added just before the closing </body> tag. For best results regarding the appearance of this code on the web page, the code should include CSS.
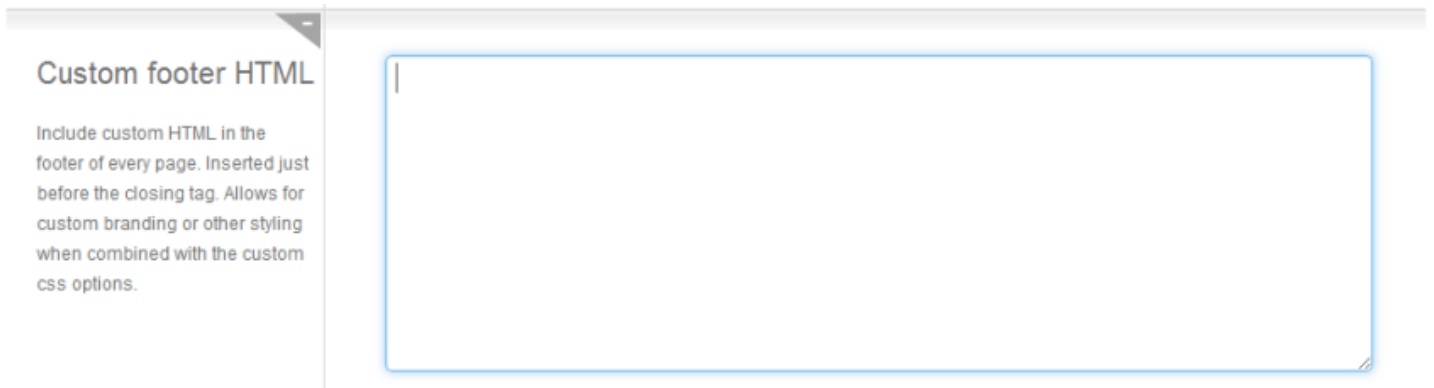
The IP Address/Host Name Groups Panel

The System Administrator can define groups of IP addresses through the IP Address Groups section of the System Settings page. IP addresses can be given individually, with wildcards, or using Netmask format. Enter only one IP address or pattern per line.

Similarly, to restrict access to an assignment to specific host names, the system administrator can define a host name pattern such that only computers that match this pattern are allowed to complete the assignment. Host names must be entered with wildcards. For example, (*.maplesoft.com).

For example, you might create groups that corresponding to computer labs at your institution.

Groups that are defined in the system settings will be available to instructors when they define the properties of assignments or tests. In the Assignment Editor, instructors can require students take an assigment from a restricted set of IP addresses or host names by selecting from a list of predefined groups set up by the system administrator or creating a custom IP address/host name restriction.

Note: To enable the host name restriction feature for instructors, the system administrator must configure Tomcat to actually do the host name lookups. This is done by adding "enableLookups" to the connector tag in the TOMCAT/conf/server.xml file. For example:

<Connector connectionTimeout="20000" port="8080" protocol="HTTP/1.1" redirectPort="8443" enableLookups="true"/>

After making this change to the connector tag, Tomcat must be restarted.

The Custom User Fields Page

In the Custom User Fields panel, you can add required fields to a user profile (for example, Grade, Program, School, Major, etc).

When these fields are added, new users are required to have these fields filled in.

Active Users Page

The Active Users Page shows those users who are currently logged into the system. Click on a User Login link to view user details.

Usage Page

The System Administrator user interface allows the system administrator to view usage data and group the results by Restart, Hour, Day, Week, Month, or Year.

Property	Default	Comments
Start Date	None	The earliest date that you want to include in your usage query. The Start Date can only be selected via the calendar.
		To use the Start Date as the beginning of the time period for the search, select use start date from the drop down list. To use the

Property	Default	Comments
		default start date for the time period, select round date to start period from the drop down list.
End Date	None	The final date that you want to include in your usage query. As with the Start Date, the End Date can only be selected via the calendar.
Include Zero Values	False (unselected)	Include those times where no one was in the system.
Show All Entries	False (unselected)	Select the Show All Entries box to see a detailed account of the system usage for the desired time period. After the Show All Entries box has been selected, re-submit the query and click on the plus ("+") sign beside the start date in the Date column of the table to see individual entries.

Note: Searches can be further refined by clicking on the "start" and "end" links in the same row as the desired individual Start Date and End Date entries in the Date column.

Results

Date	Users Avg	Max	Logins	Assignments Started	Finished	Entries		
5/25/14 12:00:00 AM EDT	1	1	2	1	0	34	start	
5/26/14 12:00:00 AM EDT	5.24	13	79	14	8	341	start	end
5/27/14 12:00:00 AM EDT	4.086	18	98	37	24	417	start	end
5/28/14 12:00:00 AM EDT	0	0	0	0	0	0		end

Log Utility Page

Note: This page will not send any information to Maplesoft Technical support unless Maple T.A. is properly connected to an SMTP email server.

For more information on logging, see *Logging (page 22)*.

Logs Folder Location

This utility will allow you to easily gather together system log files for advanced trouble-shooting.

Specify location of Tomcat's logs folder:

Name

School

To Address | support@maplesoft.com

Cc Address | (optional)

From Address

(optional)

Please specify any information you wish to include in the email (optional)

Submit

Property	Default	Comment
Specify location of Tomcat's logs folder	None	Use the full path to the logs directory. For example, on Windows: c:/Program Files/Apache/tomcat/logs or on Linux: /usr/local/apache/logs.
Name	None	Name is a required field.
School	None	School is a required field.
To Address	support@maplesoft.com	By default, the logs will be sent to Maplesoft Technical Support. However, this can be

		changed if you want to email the log files to somebody else.
Cc Address (optional)	None	Type in any email address you want to also receive a copy of this email.
Please specify any information you wish to include in the email (optional)	None	Provide additional information such as: CPU speed and amount of RAM, Operating System and versions of Tomcat, Postgres, Java and Maple T.A. (including hotfixes).

3.2 Gradebook Paging and Scrolling

Gradebook displays for administrators and instructors can be large (many thousands of rows in some cases). For this reason, Maple T.A. manages their display with a combination of paging and scrolling. This paging and scrolling is configured by properties in <Maple T.A.>/config/system.properties.

Property	Default	Comments
gradebook.tables.pagesize	50	The number of gradebook rows available for display in the browser at one time. The rows may all be displayed simultaneously or they may be vertically scrolled, with a smaller number visible simultaneously. Must be at least 10.
gradebook.tables.rowsize	25	The maximum number of columns in each row of the clickable page labels. Must be at least 10.
gradebook.tables.scrollsize	0	Defines vertical scrolling. The number of gradebook rows visible simultaneously. If 0 or >= gradebook.tables.pagesize, there is no scrolling.

3.3 Gradebook Cell Formatting Options

By default, the gradebook displays empty cells in the following way:

• For web page displays, empty cells are displayed as a dash character.

• For CSV exports, empty cells are left empty so they display as empty cells in spreadsheets.

This behavior can be customized by modifying a properties file, <Tomcat> webapps/mapleta/WEB-INF/classes/com/maplesoft/mapleta/gradebook/struts/app/ApplicationResources.properties, as follows:

Property	Default	Comments
label.decorator.number.empty	-	Can be any text.
label.decorator.csv.number.empty	<nothing>	Can be any text.

4 Monitoring and Logging

4.1 Monitoring

Maple T.A. supports monitoring of the health of key components of the system:

1. Tomcat web server

2. Postgres database server

3. Maple engines

4. X window system or Xvfb for Linux installations

Monitoring entry points allows you to perform a non-intrusive health check. You can check all components with one HTTP request or check the services individually.

To ping all components with one HTTP request, the monitoring URL is: **http://MAPLETA/ws/monitor**, where MAPLETA is the path to your Maple T.A. installation. With installation defaults, this is **http://localhost:8080/mapleta/ws/monitor**.

This URL returns a page with a single word, UP or DOWN. UP indicates that all components are up and running. DOWN indicates one or more components are not functioning properly. If DOWN is returned, you can perform separate checks to determine which service caused the problem.

To ping services separately, use the following URLs:

• Tomcat monitor — **http://MAPLETA/ws/tomcatMonitor**

• Postgres monitor — **http://MAPLETA/ws/databaseMonitor**

• Maple monitor — **http://MAPLETA/ws/mapleMonitor**

• Xvfb monitor — **http://MAPLETA/ws/xvfbMonitor**

All URLs return a page with a single word, UP or DOWN. UP indicates that the component is up and running. DOWN indicates a problem with the monitored service.

4.2 Logging

How Maple T.A. Logs

Maple T.A. logs using log4j, an Apache Software Foundation component widely used by Java applications. See **http://logging.apache.org/log4j/**.

Maple T.A. components log in such a way that different functional areas of the application can log to different files.

Configuring Logging

Logging is configured by properties in **<Maple T.A.>/config/log4j.properties**. The details are specific to the log4j product and are beyond the scope of this document. See **http://logging.apache.org/log4j/**.

The Default Log Level

The default log level set by the Maple T.A. installer is WARN for most log files. These means that log messages are limited to those indicating warnings or errors. The reason is so that the logs can be expected to contain only exceptional events, things that need action. This supports "management by exception". If you wish to see more detail in the logs, you may increase the log levels to INFO (moderate detail) or DEBUG (most detail, can be very verbose).

Contents of the Log Files

The Maple T.A. application logs to files in <Tomcat>/logs. The following table summarizes the log files and their contents:

Log File	Contents
classmanager.log	Class creation, deletion etc.

Log File	Contents
gateway.log	The taking of tests.
importer.log	Roster imports.
ldap.log	LDAP initialization, searching.
maple_router.log	Access to the pool of Maple servers.
mapleta.log	Aggregates all Maple T.A. logging – the "master log".
proctor.log	Events relating to proctor authorization of students beginning assignments.
qbeditor.log	Question Bank editing.
session.log	Login, logout.
urlrewrite.log	Information on Maple T.A.'s rewriting of URLs to facilitate control flow in the web application. This is rarely interesting.
useradmin.log	Some information on user additions and importing (overlap with importer.log).
usermanager.log	Similar to usermanager.log, will likely be merged in a future release.

Tomcat itself also logs to files in <Tomcat>/logs. The following table summarizes the log files and their contents:

Log File	Contents
catalina.log	Useful information about startup. During startup may log information about Maple T.A. that Maple T.A. itself is not able to log.
localhost.log	Messages specific to the Tomcat instance. There is rarely anything useful for Maple T.A. tracking or diagnosis in this log.
stdout.log	May log exceptions not caught by Maple T.A.
stderr.log	Like stdout.log, may log exceptions not caught by Maple T.A.
Various other Tomcat logs	Not normally useful for Maple T.A. tracking or problem diagnosis.

Dealing with the "Error listenerStart" Problem

There is a known conflict between Tomcat's logging and log4j logging by Tomcat applications. The problem's symptom is an obvious gap in log information when errors occur during certain application initialization situations.

Examples of this can happen when Maple T.A.'s access to its database is configured incorrectly and when Maple T.A. encounters security problems when Tomcat Security is enabled with a multiple-instance Tomcat web application structure, without compensating security policy configuration. In these cases you will see 'Error listenerStart" messages in Tomcat's catalina.log, and very little else.

To diagnose situations like this, it is helpful to remove the log4j jar file temporarily from <Tomcat>/webapps/mapleta/WEB-INF/lib. After that (and a Tomcat restart), Tomcat will usually log a great deal more information about the situation, helping problem diagnosis.

5 Implementing SSL

Connections to Maple T.A. are managed by your Tomcat server. Configuring SSL is largely a Tomcat operation and has little to do with T.A. itself. The Tomcat documentation contains instructions. See for example **http://tomcat.apache.org/tomcat-5.5-doc/ssl-howto.html**. In the following, we summarize the steps. Refer to the Tomcat documentation for details.

Note that it is possible to configure a Maple T.A. server to authenticate with LDAP over an SSL connection. *This feature is different from enabling SSL with Maple T.A.* It is possible to implement both features, either feature by itself, or neither. For information on authenticating via LDAP over SSL, see *Example of Maple T.A. LDAP Over SSL (page 33)*.

The following steps implement SSL for Maple T.A.:

1. Stop Tomcat.

2. Create a keystore with the certificate you want to use. You may place it anywhere, but the <Tomcat> installation directory is a convenient and reasonable place.

3. In the <Tomcat>\conf\server.xml file, make the following changes:

    a. Disable the 8080 (non-SSL) connector definition by surrounding it with XML comment tags.

    <!--
    <Connector

    port="8080"

    maxHttpHeaderSize="8192"

    maxThreads="150"

    minSpareThreads="25"

    maxSpareThreads="75"

    enableLookups="false"

    redirectPort="8443"

    acceptCount="100"

    connectionTimeout="20000"

    disableUploadTimeout="true"

    />

    -->

    b. Enable the 8443 connector definition by removing the XML comment tags that surround it.

    <Connector

    port="8443"

    maxHttpHeaderSize="8192"

    maxThreads="150"

    minSpareThreads="25"

    maxSpareThreads="75"

    enableLookups="false"

    disableUploadTimeout="true"

acceptCount="100"

scheme="https"

secure="true"

clientAuth="false"

sslProtocol="TLS"

keystoreFile="<keystore_path>"

keystorePass="<keystore_password>"

/>

Provide <keystore_path> and <keystore_password> as appropriate. <keystore_path> may be an absolute path name or it may be relative to the <Tomcat> installation folder.

4. In the <Tomcat>\webapps\mapleta\WEB-INF\web.xml file, make the following two changes:

a. In the ForwardRequest servlet section, set the port to that specified in the SSL connector (by default, 8443 – see Step 3 above):

<servlet>

...

<init-param>

<param-name>port</param-name>

<param-value>8443</param-value>

</init-param>

...

</servlet>

b. In the same section (ForwardRequest servlet), add the following:

<init-param>

<param-name>protocol</param-name>

<param-value>https</param-value>

</init-param>

5. Restart Tomcat.

6 Maple T.A. LDAP Authentication

6.1 Introduction

Maple T.A. supports two kinds of authentication:

1. Maple T.A. Authentication. Usernames and passwords are stored in Maple T.A.'s database.

2. LDAP Authentication. Usernames and passwords are stored in an LDAP server, typically administered by your institution's IT department, and used throughout the institution.

Maple T.A. supports the LDAP standard and should be able to use any LDAP server implementation. For example, we have customers using OpenLDAP®, Active Directory®, Sun™ ONE Directory Server and Novell eDirectory™. These servers are deployed on a variety of platforms including Linux and Windows Server.

6.2 Requirements for Maple T.A. LDAP Authentication

Maple T.A. LDAP requires the following for integration. If you cannot meet these requirements, you cannot integrate Maple T.A. 10 with your LDAP:

Requirement	Description
LDAP version	V3
Simple binding	Maple T.A. supports only "simple" binding. Maple T.A. does not support SASL, the alternative to simple binding.
Maple T.A. Admin in LDAP	The Maple T.A. install requests the username and password of a Maple T.A. administrator account. That username/password combination must be an entry in LDAP in one of the containers (subtrees) configured for Maple T.A. in the Maple T.A. ldap.properties file. The Maple T.A. administrator does not need any special LDAP permissions.

Search User in LDAP	Maple T.A. supports two modes – "authenticated search" and "anonymous search".
	
	1. "Authenticated search" requires configuration of a username and password for a "Search User". This may optionally be the same as the Maple T.A. Admin (see above). The Search User is required for:
	
	    a. Authentication
	
	    b. Importing student rosters
	
	    c. Adding users to Maple T.A. manually
	
	The Search User must have sufficient LDAP permissions to be able to read the values of the LDAP attributes that are mapped to Maple T.A. attributes in the configuration file. See *Configuring Maple T.A. LDAP Authentication (page 28)*.
	
	2. "Anonymous search" requires no username or password. In this case, an anonymous bind must have sufficient LDAP permissions to locate users and read the values of the LDAP attributes that are mapped to Maple T.A. attributes in the configuration file. See *Configuring Maple T.A. LDAP Authentication (page 28)*.
	
	These modes and any search user credentials may be specified at the subtree (container) level. For example, you may have different search users in different subtrees.

Maple T.A. login attribute	You must select an LDAP attribute whose value is the text typed in by Maple T.A. users for "User login" in the login dialog. This is often uid for Linux LDAP servers, or sAMAccountName for Active Directory. There are other options, depending on the LDAP repository structure. The overriding constraint is that the login attribute must uniquely identify a user. login attributes may be specified at the subtree (container) level in which case the uniqueness constraint applies only to the subtree.

Requirement	Description
Maple T.A. username attribute in LDAP	The repository must provide a user attribute that uniquely identifies users within the space of Maple T.A. users. In practice, uid or cn can usually be used for this. It is not essential that this attribute be the "login attribute" for Maple T.A., or the rdn for LDAP, although that is frequently the case. Be aware that the value of this attribute is displayed in columns titled "Login" throughout Maple T.A.
Maple T.A. Last Name attribute in LDAP	The repository must provide a last name attribute. In practice, sn can usually be used for this. If a last name attribute is not available, last name can be derived from a fullname attribute using an attribute parser.
Maple T.A. First Name attribute in LDAP	The repository must provide a first name attribute. In practice, givenName can usually be used for this. If a first name attribute is not available, first name can be derived from a fullname attribute using an attribute parser.

Be aware that, besides username, last name, and given name, Maple T.A. can also use the following optional items of information from LDAP, to describe users. These items appear in various tables and reports:

• Middle initial

• Email address (non-null values must be unique)

• Student Id (non-null values must be unique)

6.3 Specifying Maple T.A. LDAP Authentication

You select the authentication mode during Maple T.A. installation.

6.4 Changing the Authentication Mode After Installation

It is possible to change the mode after installation, although the need for this is extremely rare. The mode is changed by modifying the Maple T.A. database directly with a small package of SQL commands. See *Using PostgreSQL Tools (page 39)* for information on how to execute SQL commands.

Changing from LDAP to Maple T.A.:

update system_properties set value = 'DBSecurityServiceImpl' where key = 'AuthenticationBean';

update system_properties set value = '' where key = 'ProfileImporterBean';

update system_properties set value = 'false' where key = 'LDAP_ENABLED';

Changing from Maple T.A. to LDAP:

update system_properties set value = 'LDAPSecurityServiceImpl' where key = 'AuthenticationBean';

update system_properties set value = 'LDAPProfileImporterServiceProviderImpl' where key = 'ProfileImporterBean';

update system_properties set value = 'true' where key = 'LDAP_ENABLED';

Be aware that if you change authentication from LDAP to Maple T.A., the passwords will be unknown for all users except the admin specified during installation. In particular, the passwords will not be those that are in the LDAP repository. You can do a batch update of user passwords with the roster import facility.

6.5 Diagnostics

Information pertinent to LDAP appears in the session and LDAP logs. These files are:

1. <Tomcat>/logs/session.log

2. <Tomcat>/logs/ldap.log

For problem diagnosis, session and LDAP logging should be set to DEBUG in the <Maple T.A.>/config/log4j.properties file.

The relevant lines are

log4j.category.com.maplesoft.mapleta.session=DEBUG, R, session

log4j.category.com.maplesoft.mapleta.ldap=DEBUG, R, ldap

Note that you must restart Tomcat in order for log4j.properties changes to take effect.

6.6 The Maple T.A. LDAP Authentication Model

Maple T.A. LDAP does three things:

1. Searches for users in LDAP

2. Authenticates users in LDAP

3. Maps LDAP user attributes to Maple T.A. user profile attributes

Searching for users is always done under a "search user", which may be (rarely) "anonymous". Searching for users is organized by subtree. You configure subtrees and rules for searching them. The configuration refers to subtrees and their rules as "Containers". It is possible to have one container, or several. Containers can inherit from a common base container, simplifying the configuration description.

Authenticating a user in LDAP is always done by "binding" the user directly to LDAP, using the user's password.

A Maple T.A. user has a profile comprising seven attributes. Each attribute has an internal name and a display name. The internal name is used in Maple T.A.'s database schema and in configuring Maple T.A. LDAP. The display name is used in the Maple T.A. web application's pages. The following table summarizes:

Internal Name	Display Name	Comment
uid	Login	Required, must be unique within Maple T.A.
givenName	First Name	Required.
sn	Last Name	Required.
mi	MI	Optional.
cn	cn	Optional. Not exposed in any Maple T.A. web pages currently.
studentID	Student Id	Optional. Specified Student Id's must be unique within Maple T.A.
email	Email Address	Optional. Specified Email Addresses must be unique within Maple T.A.

6.7 Configuring Maple T.A. LDAP Authentication

You configure Maple T.A. LDAP with a properties file, ldap.properties, found in the config folder of your Maple T.A. installation.

Lines in this properties file have the form "property" = "value". Empty lines and lines starting with the "#" character (comment lines) are skipped.

The following table summarizes the properties:

Property	Required	Value
Global.autorefresh	No	'true' or 'false'. The default is 'false'. If 'true', Maple T.A. will detect changes in the ldap.properties file and reload them immediately. If 'false', you must restart Maple T.A. (Tomcat) for changes to take effect.
Global.server.url	Yes	examples: 1. ldap://myrepository.myu.edu 2. ldaps://myrepository.myu.edu (for SSL)
Global.server.ssl.trustStore	No	Required for SSL. File path. Under Windows, any '\' (backslash) characters in the path must be doubled, e.g. C:\\Documents and Settings\\... Alternatively under Windows,

Property	Required	Value
		you may use the Unix-style '/' (forward slash) as the path component separator, e.g. C:/Documents and Settings/...
Global.server.ssl.trustStorePassword	No	Required for SSL. Text string.
Global.server.ssl.keyStore	No	For SSL, optional (normally not necessary but this depends on your LDAP server's SSL configuration). Under Windows, same rules apply for path components separators as for trustStore.
Global.server.ssl.keyStorePassword	No	For SSL, required only if keyStore is required.
Container<id>.Search.User	Yes	ldap full dn or "anonymous" (without quotes).
Container<id>.Search.Password	Yes/No	Password for Container<id>.Search.User. Required if Search.User is not "anonymous".
Container<id>.base	Yes	dn suffix used as a base for searches within the container. Example dc=myu, dc=edu. Containers other than ContainerCommon concatenate their bases with ContainerCommon's base to form the search base. For example: ContainerCommon.base = dc=myu,dc=edu Container1.base = ou=Students Then searches in Container1 will be based at ou=Students,dc=myu,dc=edu.
Container<id>.login.attribute	No	The ldap attribute that is used to login with. Defaults to the value of Container<id>.ta.uid.
Container<id>.ta.uid	Yes	ldap.xxx where xxx us the ldap attribute to be mapped to Maple T.A.'s uid (login) attribute.
Container<id>.ta.givenName	Yes	ldap.xxx where xxx us the ldap attribute to be mapped to Maple T.A.'s givenName (first name) attribute.
Container<id>.ta.sn	Yes	ldap.xxx where xxx us the ldap attribute to be mapped to Maple T.A.'s sn (last name) attribute.
Container<id>.ta.mi	No	ldap.xxx where xxx us the ldap attribute to be mapped to Maple T.A.'s mi (middle initial) attribute.
Container<id>.ta.cn	No	ldap.xxx where xxx us the ldap attribute to be mapped to Maple T.A.'s cn (common name) attribute.
Container<id>.studentID	No	ldap.xxx where xxx us the ldap attribute to be mapped to Maple T.A.'s studentID attribute.
Container<id>.email	No	ldap.xxx where xxx us the ldap attribute to be mapped to Maple T.A.'s email attribute.
Container<id>.ta.uid.parser	No	A regular expression enclosed with "/" characters and with a single capturing group. The regular expression is applied to the ldap attribute that is mapped to the ta login attribute. The ta login attribute receives the value of the capturing group.
Container<id>.ta.givenName.parser	No	Analogous to ta.uid.parser, but for the givenName attribute.
Container<id>.ta.sn.parser	No	Analogous to ta. uid.parser, but for the sn attribute.
Container<id>.ta.mi.parser	No	Analogous to ta. uid.parser, but for the mi attribute.
Container<id>.ta.cn.parser	No	Analogous to ta. uid.parser, but for the cn attribute.
Container<id>.studentID.parser	No	Analogous to ta. uid.parser, but for the studentID attribute.
Container<id>.email.parser	No	Analogous to ta. uid.parser, but for the emailAddress attribute.

Notes:

1. <id> may be the text "Common" or a number 1, 2, 3, …. For example, the following are valid:

    a. ContainerCommon.base

    b. Container1.base

2. ContainerCommon establishes values that are inherited by Container1, Container2, … This is best understood through the examples in the following section.

3. Container1, Container2, etc. cannot have gaps. That is, if you have defined Container<n>, you must have defined Container<n-1>

4. You may omit the Container<n> definitions altogether. In that case, Maple T.A. creates a default Container1 that uses the values from ContainerCommon.

5. All property names are case-insensitive. That is, upper and lower case characters are treated as the same. However, property values are case sensitive.

6. Maple T.A. attributes are identified, for mapping, using their internal names, not their display names (see *The Maple T.A. LDAP Authentication Model (page 28)*).

7. The Maple T.A. attribute cn ("common name") is not used in any Maple T.A. reports. It is not useful to map it unless you anticipate custom querying of The Maple T.A. database for which cn might be needed.

6.8 Maple T.A. LDAP Configuration Examples

Simple OpenLDAP Single-Container

This example is not specific to OpenLDAP but it assumes the default schema installed by OpenLDAP.

ldap.properties:

Global.server.url = ldap://myrepository.myu.edu

ContainerCommon.Search.User = uid=mapleta, ou=Administrators, dc=myu, dc= edu

ContainerCommon.Search.Password = secret

ContainerCommon.base = ou=People, dc= myu, dc = edu

ContainerCommon.ta.uid = ldap.uid

ContainerCommon.ta.givenName = ldap.givenName

ContainerCommon.ta.sn = ldap.sn

Operation

Logging in

1. Student enters jdoe and his password in the Maple T.A. login dialog.

2. Maple T.A. logs in to LDAP as dn=[uid=mapleta, ou=Administrators, dc=myu, dc= edu].

3. Maple T.A. searches for uid=jdoe in the subtree ou=People, dc=myu, dc =edu.

    a. If not found, login is rejected.

    b. If found, Maple T.A. attempts to log in as dn=[uid=jdoe, ou=People, dc=myu, dc =edu] with the password entered by jdoe in the Maple T.A. login dialog.

        i. If not successful, login is rejected.

        ii. If successful, login is accepted.

            A. If there is a record with login jdoe in the Maple T.A. database, Maple T.A. reads givenName, sn and updates firstName, lastName respectively, if they have changed since the last login (Maple T.A. saves these values in its database for dispaly purposes).

            B. If there is no record with login jdoe in the Maple T.A. database, Maple T.A. may or may not add the user to its database, depending on the configuration options for "self registration".

Roster Import

1. Maple T.A. logs in to LDAP as dn=[uid=mapleta, ou=Administrators, dc=myu, dc= edu].

2. For each line in the import file:

    a. Maple T.A. searches for uid=xxx in the subtree ou=People, dc=myu, dc =edu where xxx is value of the Login column in the import line.

        i. If not found, the line is rejected.

        ii. If found, Maple T.A. searches its database for login xxx. If found, the user record is updated (including refreshing of firstName and lastName) according to the import operation type. If not found, the user may be added to the database, depending on the import operation type.

Add User

This dialog, under Administer Users, displays a dropdown of searchable attributes. These are precisely the "friendly names" of ta attributes that are mapped from LDAP attributes. The administrator selects an attribute to search on, and a value. For example, "Last Name" and "Do".

1. Maple T.A. logs in to LDAP as dn=[uid=mapleta, ou=Administrators, dc=myu, dc= edu].

2. Maple T.A. searches for sn=Do* in the subtree ou=People, dc=myu, dc =edu.

3. Maple T.A. displays the users that were found and you can select any of them that you wish to add to Maple T.A.

Simple OpenLDAP Single-Container with Anonymous Search

This example is almost identical to *Simple OpenLDAP Single-Container (page 30)*. The only difference is in the choice of the anonymous search user.

ldap.properties:

Global.server.url = ldap://myrepository.myu.edu

ContainerCommon.Search.User = anonymous

ContainerCommon.base = ou=People, dc= myu, dc = edu

ContainerCommon.ta.uid = ldap.uid

ContainerCommon.ta.givenName = ldap.givenName

ContainerCommon.ta.sn = ldap.sn

Operation

Same as the operation of *Simple OpenLDAP Single-Container (page 30)*".

Simple OpenLDAP Multi-Container

This example is not specific to OpenLDAP but it assumes the default schema installed by OpenLDAP. This example assumes the "People" LDAP node is divided into two subtrees – "Students" and "Staff".

ldap.properties:

Global.server.url = ldap://myrepository.myu.edu

ContainerCommon.Search.User = uid=mapleta, ou=Administrators, dc=myu, dc= edu

ContainerCommon.Search.Password = secret

ContainerCommon.base = ou=People, dc= myu, dc = edu

ContainerCommon.ta.uid = ldap.uid

ContainerCommon.ta.givenName = ldap.givenName

ContainerCommon.ta.sn = ldap.sn

Container1.base = ou=Students

Container2.base= ou=Staff

Operation

Logging in

This is identical to the previous single-container example, except that the search in step 3 above is tried first in the subtree ou=Students, ou=People, dc= myu, dc=edu and then, if that fails, in the subtree ou=Staff, ou=People, dc= myu, dc=edu.

Roster Import and Add User

Again, this is identical to the previous single-container example, except that the search is tried in both subtrees if necessary, as for login.

Active Directory

Active Directory by default requires that cn be used for authentication. By default this is the full name – first and last ("John Doe") and not usually what people want to use for logging in. The login preference is often userPrincipalName ("jdoe@myu.edu") or sAMAccountName ("jdoe").

Here is a single-container example that has users login with userPrincipalName, but stores sAMAccountName in the Maple T.A. LDAP login attribute.

ldap.properties:

Global.server.url = ldap://myrepository.myu.edu

ContainerCommon.Search.User = cn=Maple TA, ou=Administrators, dc=myu, dc= edu

ContainerCommon.Search.Password = secret

ContainerCommon.base = ou=People, dc= myu, dc = edu

ContainerCommon.login.attribute = userPrincipalName

ContainerCommon.ta.uid = ldap.sAMAccountName

ContainerCommon.ta.givenName = ldap.givenName

ContainerCommon.ta.sn = ldap.sn

Operation

Logging in

1. Student enters jdoe@myu.edu and his password in the Maple T.A. login dialog.

2. Maple T.A. logs in to LDAP as dn=[cn=mapleta, ou=Administrators, dc=myu, dc= edu].

3. Maple T.A. searches for userPrincipalName=jdoe@myu.edu in the subtree ou=People, dc=myu, dc =edu.

4. If not found, login is rejected.

5. If found, Maple T.A. knows cn as a byproduct of the search and attempts to log in as dn=[cn=John Doe, ou=People, dc=myu, dc =edu] with the password entered by jdoe@myu.edu in the Maple T.A. login dialog.

6. If not successful, login is rejected.

7. If successful, login is accepted.

8. Maple T.A. reads the attribute SaMAccountName=jdoe. If there is a record with login jdoe in the Maple T.A. database, Maple T.A. reads givenName, sn from LDAP and updates firstName, lastName respectively, if they have changed since the last login (Maple T.A. saves these values in its database for display purposes).

9. If there is no record with login jdoe in the Maple T.A. database, Maple T.A. may or may not add the user to its database, depending on the configuration options for "self registration".

Roster Import and Add User

The search and identification of Maple T.A. users is done as in the Logging In case.

Example of an Attribute Parser

Suppose there is an employeeNumber in LDAP that matches what you want to use for a student id, except for a prefix "edu:" that you would like to discard. So, for example if employeeNumber is "edu:12345", you want to use "12345" as the student id. This is a simple example of what can be done with an attribute parser.

ldap.properties:

Global.server.url = ldap://myrepository.myu.edu

ContainerCommon.Search.User = uid=mapleta, ou=Administrators, dc=myu, dc= edu

ContainerCommon.Search.Password = secret

ContainerCommon.base = ou=People, dc= myu, dc = edu

ContainerCommon.ta.uid = ldap.uid

ContainerCommon.ta.givenName = ldap.givenName

ContainerCommon.ta.sn = ldap.sn

ContainerCommon.ta.studentID = ldap.employeeNumber

ContainerCommon.ta.studentID.parser = /edu:(.*)/

Operation

When Maple T.A. reads the employeeNumber in order to compute ta.studentID, it matches against the regular expression specified by ta.studentID.parser and extracts the value of the capturing group. For employeeNumber = edu:12345, the value captured by the parser is 12345, as desired.

Example of Maple T.A. LDAP Over SSL

Imagine the server in the example *Simple OpenLDAP Single-Container (page 30)* requires a secure connection using SSL.

You need to import the server's certificates to the trustStore file on the Maple T.A. server (or accessible to the Maple T.A. server). You may optionally need to import certain related information into a keyStore file. The details of this are part of the realm of SSL and beyond the scope of this document.

To enable LDAP over SSL, you must:

1. Specify SSL in the server URL by using the protocol prefix ldaps:// (rather than ldap://).

2. Identify the location of the trustStore and optionally the keyStore, and their passwords.

ldap.properties:

Global.server.url = ldaps://myrepository.myu.edu

Global.server.ssl.trustStore = /usr/share/MapleT.A.4.0/ssl/ta.jks

Global.server.ssl.trustStorePassword = sslsecret

ContainerCommon.Search.User = uid=mapleta, ou=Administrators, dc=myu, dc= edu

ContainerCommon.Search.Password = secret

ContainerCommon.base = ou=People, dc= myu, dc = edu

ContainerCommon.ta.uid = ldap.uid

ContainerCommon.ta.givenName = ldap.givenName

ContainerCommon.ta.sn = ldap.sn

Operation

Maple T.A. operation in this configuration is identical to that of *Simple OpenLDAP Single-Container (page 30)*.

7 Applying Hotfixes

Maplesoft has a systematic approach to providing Maple T.A. hotfixes to customers for urgent problems.

Hotfixes are very small patches that fix specific highly focused problems. Hotfixes are installed using Maple T.A.'s "Hotfix Installer" tool, which is bundled in each hotfix package (it is very small). The Hotfix Installer applies fixes by modifying files in the deployed Maple T.A. web application, that is, by modifying files in <Tomcat>/webapps/mapleta.

Hotfixes are published on Maplesoft's web site under **Support>Download Product Updates**. See **http://www.maplesoft.com/support/downloads/index.aspx**.

Hotfixes are distributed as compressed archives–".zip" for Windows and ".tar.gz" for Linux . Installation of a hotfix is accomplished by unzipping the contents of the archive and running a .bat or .sh. Normally installation is completely automatic. Documentation on how to install hotfixes is included with every hotfix package.

Notes:

1. The Hotfix Installer is a Java program. It relies on the Java environment that is a prerequisite for T.A. Your Java bin directory must be in your system path. See *Installing Java on the Server (page 3)*.

2. The Hotfix Installer verifies the version number of the Maple T.A. installation. Hotfixes are version-specific.

3. The Hotfix Installer automatically backs up the files it changes, before it changes them. It is possible to revert the changes afterwards by running the install script with the "-revert" parameter.

4. The Hotfix Installer detects "fix already installed" or "fix partially installed". It will not install the fix in these situations, unless you run the install script with the "-force" parameter.

5. The Hotfix Installer logs details of its actions in a log subfolder of the hotfix folder. It also logs a one-line record of the install in <Maple T.A.>/hotfixlog/hotfixes.log. It is always possible to check what hotfixes you have installed by looking at the hotfixes.log file.

6. The Hotfix Installer normally finds <Tomcat> and <Maple T.A.> automatically. If <Tomcat> is in an unusual place, the installer may not find it. In that case, you can provide the path to the installer directly by modifying a text file (hotfix.def) in the hotfix folder.

8 Sending Diagnostic System Snapshots to Maplesoft

In very unusual problem cases, resolution is sometimes helped by providing the Maplesoft Maple T.A. team with a snapshot of your installation so they can reproduce the problem. This situation is rare but it has arisen in the past.

In these cases, Maplesoft Technical Support will ask you to send the following:

1. A zip of your <Maple T.A.>/records folder.

2. A compressed backup of your PostgreSQL database. You can take this backup with pgadmin or psql (see *Using PostgreSQL Tools (page 39)*).

3. A zip of your <PostgreSQL>/data/pg_log folder.

Maplesoft commits to keeping your data confidential and to destroying it immediately upon completion of the analysis. Be aware that if you have chosen to encrypt your passwords (the installation default), it will not be possible for Maplesoft staff (or anyone else) to see your passwords.

Appendix A Summary of Maple T.A. Configuration Files

This appendix summarizes the various places in which Maple T.A. configuration information can be entered. Configuration details are described elsewhere in this document.

File	Controls
<Maple T.A.>/config/ldap.properties	LDAP interface.
<Maple T.A.>/config/log4j.properties	Logging.
<Maple T.A.>/config/system.properties	Maple server pool, MapleNet interface, gradebook UI.
<Tomcat>/webapps/mapleta/METAINF/context.xml	Access to Maple T.A.'s database. Note that Tomcat may further deploy this file, as mapleta.xml, to the instance subfolder of <Tomcat>/conf. For example, for a "localhost" instance this folder is <Tomcat>/conf/Catalina/localhost.

Appendix B Using the Tomcat Manager Facility

B.1 Basics

The Tomcat Manager is described in the Tomcat User's Guide; for example, **http://tomcat.apache.org/tomcat-6.0-doc/manager-howto.html**.

In normal Maple T.A. operations, you will never need to use the Tomcat Manager. The cases where it might be useful include:

1. Reinstalling Maple T.A. (see *Reinstalling Maple T.A. (page 8)*). This would be a rare occurrence.

2. Deploying a beta version of Maple T.A. from Maplesoft. This would be a rare occurrence.

3. Viewing the number of sessions active in Maple T.A. This could be something you might do from time to time.

The Tomcat Manager can be accessed from a link in the upper left of the Tomcat home page. You will be asked for credentials. See the following sections on Installation for instructions on how to set up credentials.

B.2 Installation

Windows

The Tomcat Manager is installed automatically by the standard Tomcat "Windows Service Installer". During the installation, you will be asked for an administrative username andpassword. These credentials will allow you to access the Tomcat Manager.

Linux

The Tomcat Manager is not installed by most Linux Tomcat packages. You need to install an additional package, usually named something like "tomcat-admin". When you do install the Tomcat Manager, the installer normally will not ask you for credentials and will not provide for any credentials. You will not be able to use the Tomcat Manager application until you create credentials. By default, the credentials must be created in <Tomcat>/conf/tomcat-users.xml. See **http://tomcat.apache.org/tomcat-6.0-doc/manager-howto.html**. The following is a simple tomcat-users.xml that will allow username "mapleta" to access the Tomcat Manger with password "secret":

<?xml version='1.0' encoding='utf-8'?>

<tomcat-users>

<role rolename="manager"/>

<role rolename="admin"/>

<user username="mapleta" password="secret" roles="admin,manager"/>

</tomcat-users>

Appendix C Using PostgreSQL Tools

It is sometimes necessary for a System Administrator to "patch" the Maple T.A. database using SQL. This is rare, but it can happen in very unusual situations. Therefore, it is helpful to have some understanding of the tools that are available.

1. The most important tool is a basic command-line SQL tool called psql, provided by PostgreSQL in every distribution. From the PostgreSQL documentation:

   psql is a terminal-based front-end to PostgreSQL. It enables you to type in queries interactively, issue them to PostgreSQL, and see the query results. Alternatively, input can be from a file. In addition, it provides a number of meta-commands and various shell-like features to facilitate writing scripts and automating a wide variety of tasks.

   You can read about psql in the PostgreSQL documentation, for example **http://www.postgresql.org/docs/8.3/interactive/app-psql.html**.

2. A useful alternative to psql, in some cases, is the GUI tool pgadmin (also called pgAdmin III). pgadmin does not come from PostgreSQL. Rather, it comes from a separate open source project – see **http://www.pgadmin.org**. pgadmin requires a windowing environment so is not suitable for some servers. However, in cases where a windowing environment is available, pgadmin may be easier to use for some System Administrators. You can read about pgadmin in its online documentation, for example **http://www.pgadmin.org/docs/1.8/index.html**.