



LISTSERV® Maestro 4.0 Administrator's Manual

www.lsoft.com



Information in this document is subject to change without notice. Companies, names, and data used in examples herein are fictitious unless otherwise noted. L-Soft does not endorse or approve the use of any of the product names or trademarks appearing in this document.

Permission is granted to copy this document, at no charge and in its entirety, if the copies are not used for commercial advantage, the source is cited, and the present copyright notice is included in all copies. Recipients of such copies are equally bound to abide by the present conditions. Prior written permission is required for any commercial use of this document, in whole or in part, and for any partial reproduction of the contents of this document exceeding 50 lines of up to 80 characters, or equivalent. The title page, table of contents, and index, if any, are not considered to be part of the document for the purposes of this copyright notice, and can be freely removed if present.

Copyright © 2010 L-Soft Sweden AB
All Rights Reserved Worldwide.

LISTSERV is a registered trademark licensed to L-Soft Sweden AB and L-Soft international, Inc. ListPlex, CataList, and EASE are service marks of L-Soft international, Inc.

The Open Group, Motif, OSF/1 UNIX and the "X" device are registered trademarks of The Open Group in the United State and other countries.

Digital, Alpha AXP, AXP, Digital UNIX, OpenVMS, HP, and HP-UX are trademarks of Hewlett-Packard Company in the United States and other countries.

Microsoft, Windows, Windows 2000, Windows XP, and Windows NT are registered trademarks of Microsoft Corporation in the United States and other countries.

Sun, Solaris, SunOS, and PMDF are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

IRIX is a registered trademark of Silicon Graphics, Inc. in the United States and other countries. Linux is a registered trademark of Linus Torvalds.

Intel and Pentium are registered trademarks of Intel Corporation.

All other trademarks, both marked and not marked, are the property of their respective owners.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

Some portions licensed from IBM are available at <http://oss.software.ibm.com/icu4j/>

This product includes code licensed from RSA Security, Inc.

Manuals are available in PDF and HTML format at <http://www.lsoft.com/manuals/index.html>

L-Soft invites comment on its manual. Please feel free to send your comments by email to manuals@lsoft.com

Last Updated: December 15, 2010

Table of Contents

Understanding LISTSERV Maestro

Preface - About This Manual	xiii
What's New in LISTSERV® Maestro 4.0	xv
Section 1 Introduction to Maestro Administration	1
1.1 Maestro User Roles	2
1.2 Remote Administration Access	3
1.2.1 Remote Log Access	3
1.2.2 Remote Version Query	5
1.3 Client System Requirements	5
1.4 Accessing the Hub Administrator Interface	6
1.4.1 Switching Between Interfaces	7
1.5 Understanding the Hub Administrator Interface	8
1.5.1 The Toolbar	8
Section 2 Configuring LISTSERV Maestro for First Use	13
Section 3 Changing the Administrator Password	15
Section 4 LISTSERV and LISTSERV Maestro	17
4.1 Preparing LISTSERV for LISTSERV Maestro	17
4.1.1 Preparing LISTSERV to Process DISTRIBUTE Jobs from LISTSERV Maestro	17
4.1.2 Preparing LISTSERV to Allow Maestro Hosted LISTSERV Lists	19
4.1.3 Preparing LISTSERV for Database Access	21
4.2 Preparing LISTSERV Maestro to Send DISTRIBUTE Jobs to LISTSERV	25
4.2.1 Specifying the LISTSERV Host with Different Internal and External Names	26
4.2.2 Specifying a Separate LISTSERV Instance for Processing Bounces	27
4.3 Using Existing Lists with LISTSERV Maestro	28
Section 5 Settings for the Maestro User Interface	29
5.1 Application Settings	29
5.1.1 General Administration	29
5.1.2 Database Plugins	31
5.1.3 System Database Connection	32
5.1.4 User Rights	32
5.1.5 LISTSERV Web Interface Access	33
5.2 Application Default Settings	33
5.2.1 Default URL Settings	34
5.2.2 Default LISTSERV Connection	36
5.2.3 Default Size Limits	39
5.2.4 Default Content Restrictions	39
5.2.5 Default Recipients Restrictions	40
5.2.6 Default Tracking Restrictions	42
5.2.7 Default Hosted Data Settings	42

5.2.8 Default DomainKeys Settings	45
5.2.9 Default Auto-Archive Settings	46
5.3 Setting the Default Tracking URL	46
5.4 Automatically Logging into the Maestro User Interface	48
Section 6 Administrative Policies	51
6.1 Configuring Backups	51
6.2 Runtime Administration and System Shutdown	51
6.3 User Restrictions	53
6.4 Administrative Email Notifications	54
6.4.1 Testing Email Notifications	56
6.5 Refreshing the Subscriber Page Translations	57
Working with Accounts and Identities	
Section 7 Creating and Editing User Accounts and Identities	59
7.1 Creating a New User Account	60
7.2 Creating a New Identity	61
7.3 Managing User Rights	62
7.4 Editing Account Information and Assigning Single User Settings	63
7.4.1 Editing General User Settings	64
7.4.2 Editing Component Specific Settings for Single and Group Users	64
7.5 Deleting a User Account, Identity, or Group	77
Section 8 Special Administrative User Account	81
8.1 The Toolbar	81
8.1.1 Behavioral Changes to the Toolbar	90
8.2 The Dashboard	90
8.2.1 Customizing the Dashboard	92
8.2.2 System Metrics	93
8.2.3 System Metrics Reports	94
8.2.4 Managing Dashboard Reports	96
8.2.5 Creating Multiple Dashboards	98
8.2.6 Switching Dashboards	99
8.2.7 Deleting Dashboards	100
8.3 Sorting and Filtering Jobs	100
8.4 Archiving Delivered/Completed Jobs	101
8.4.1 Auto-Archiving	102
8.5 Importing Archived Jobs	102
8.6 Changing Job and Report Ownership	103
8.7 Changing Sender Profile, Drop-In, and Content Template Ownership	107
8.8 Changing Ownership of the Entire Recipient Warehouse	108
8.8.1 Changing Target Group Ownership	110
8.8.2 Changing Ownership of Recipient Datasets and Lookup Tables	112
Working with Databases	
Section 9 Defining External Database Connections	115
9.1 Available Database Plugins	116
9.1.1 The IBM DB2 V8.2 Thin Driver Database Plugin	116

9.1.2 The IBM DB2 V7.2 Native Driver Database Plugin	117
9.1.3 The MySQL ConnectorJ Driver Database Plugin	117
9.1.4 The MySQL L-Soft Driver Database Plugin	118
9.1.5 The Oracle 8i, 9i, and 10g Thin Driver Database Plugin	118
9.1.6 The SQL Server jTDS Driver Database Plugin	119
9.1.7 The SQL Server Microsoft Driver Database Plugin	119
9.1.8 The SQL Server i-net SPRINTA Driver Database Plugin	120
9.1.9 The ODBC Driver Database Plugin	120
9.2 Registering a Database Plugin	122
Section 10 The System Database	123
10.1 Configuring the External System Database	124
10.2 Preparing the System Database	126
10.2.1 General System Database Preparation	126
10.2.2 Preparing SQL Server as the System Database	127
10.2.3 Preparing Oracle as the System Database	127
10.2.4 Preparing DB2 as the System Database	128
10.2.5 Preparing MySQL as the System Database	129
10.3 General Optimization Hints for the System Database	130
10.4 Removing and Adding the Internal Database	130
Section 11 Saving and Restoring a Backup	131
11.1 Configuring the Backup Time	131
11.2 Configuring External Post-Backup Processes	132
11.3 Configuring the Backup Location	133
11.4 Configuring the Backup History	134
11.5 Saving a Backup to an External Medium	134
11.6 Identifying the Backup: The Backup ID	135
11.7 Restoring a Backup	135
Section 12 Using a Test-Bed Backup	141
12.1 Creating a Test-Bed Backup on the Original System	143
12.2 Restoring a Test-Bed Backup into the Test System	144
Advanced Features	
Section 13 Maestro Logs	145
13.1 Remote Log Access	145
13.2 Subscriber Activity Change Log	146
Section 14 Using Non-Standard Ports	149
14.1 Ports Used by LISTSERV Maestro	149
14.1.1 Ports used by the Administration Hub	149
14.1.2 Ports used by the Maestro User Interface	149
14.1.3 Ports used by Maestro Tracker	149
14.2 Configuring Port Usage	150
14.2.1 Configuring the HTTP Port	150
14.2.2 Configuring the HTTPS Port	153
14.2.3 Configuring the Internal Communication Port	153
14.2.4 Configuring the Tracker Communications Port	154

14.2.5 Configuring the Internal Database Connection Port	154
14.2.6 Configuring the Application Server Shutdown Port	155
Section 15 Defining the IP Addresses and HTTP Ports	157
Section 16 Installing Behind a Firewall	159
Section 17 Restricting Access to Components	163
17.1 IP Address Restrictions	163
17.2 Disallowing Concurrent Access with the Same User Account	165
17.3 Securing Access Against Dictionary Attacks	168
17.3.1 Securing the Administration Hub	168
17.3.2 Securing the LISTSERV Maestro User Interface	168
Section 18 Securing Access with SSL	169
18.1 Introduction to Secure Communication	169
18.2 Which Components Should Be Secured?	172
18.3 Obtaining and Installing a Server Certificate	173
18.3.1 Securing the Trusted Root Certificate Keystore	173
18.3.2 Creating an Unsigned Server Certificate	174
18.3.3 Performing a Certificate Signing Request (CSR)	175
18.3.4 Installing the Signed Server Certificate	176
18.3.5 Installing a Trusted Root Certificate	177
18.3.6 Making LISTSERV Maestro Aware of the Server Certificate	178
18.3.7 Securing a Server with Multiple Host Names	180
18.3.8 Mixing SSL and Non-SSL Access on One Server (Mixed Mode)	181
18.3.9 Supported SSL Cipher Suites	183
Section 19 Tracking and Recipient Profiles	185
Section 20 Editing LISTSERV Maestro INI Files	189
20.1 Maestro User Interface INI-File Entries	190
20.2 Administration Hub INI-File Entries	195
20.3 Maestro Tracker INI-File Entries	196
20.4 Tomcat INI-File Entries	197
20.4.1 Basic Tomcat Configuration Parameters	197
20.4.2 Advanced Tomcat Configuration Parameters	198
Section 21 Authenticating Message Origin with DomainKeys Signatures	201
Section 22 LISTSERV & LISTSERV Maestro Integration	205
22.1 Defining the LISTSERV and LISTSERV Maestro Interface Links	205
22.2 Enabling Single Sign-On	209
22.3 Linking the Membership Area and the Subscriber's Corner	213
Section 23 Moving a LISTSERV Maestro Installation	217
23.1 Server Replacement	218
23.2 Moving a LISTSERV Maestro Installation	220
23.3 Preparation for Replacing/Moving: Saving a Backup	223

Section 24 Distributed Components	225
24.1 Fresh Installation with Distributed Components	226
24.2 Moving Components to another Server	226
24.2.1 Moving the Maestro User Interface Component to Another Server.....	226
24.2.2 Moving the Administration Hub Component to Another Server	229
24.2.3 Moving the Maestro Tracker Component to Another Server	230
24.2.4 Moving the Database External Component to Another Server	232
24.3 Server Name Aliases and Proxies	232
24.3.1 Configuring LISTSERV Maestro Components with Server Name Aliases or Proxies	234
Section 25 LISTSERV Maestro in Evaluation Mode	237
Section 26 Adding Content to the Tomcat Server	239
26.1 Adding the LISTSERV Web Interface to the Tomcat Server	239
26.2 Adding Custom Content to the Tomcat Server	240
26.2.1 Adding Content as a New Context	240
26.2.2 Defining the Default Context	242
26.3 Enabling Access Logging for Added Content	244
26.3.1 Enabling Access Logging for WA	244
26.3.2 Enabling Access Logging for Custom Content	245
26.4 ListPlex and the Tomcat Server	246
Section 27 Using International Character Sets	253
27.1 Defining the Default Mail Charset	253
27.2 Allowing or Disallowing Bi-Directional Character Sets	254
Section 28 LISTSERV Maestro Standard Default Ports	257
Section 29 Updating Maestro's HTML Upload Applet	259
29.1 Sun Java-Plugin	259
29.2 Microsoft Java Virtual Machine	260
Section 30 LISTSERV Maestro Lite	263
30.1 Current Lite-Mode Restrictions	264
30.1.1 User-Feature Restrictions	264
30.1.2 Admin-Feature Restrictions	267
Section 31 Emergency Admin HUB Access	269
Section 32 Sharing a Server with IIS	271
32.1 Configuring LISTSERV Maestro	271
32.2 Configuring IIS	272
32.2.1 Configuring IIS 6	273
32.2.2 Configuring IIS 7	274
32.3 Completing the Configuration	274
Glossary and Index	
Glossary of Terms	277
Index	285

List of Figures

Accessing LISTSERV Maestro	6
LISTSERV Maestro Login Screen	7
Switching Between Interfaces	7
The Toolbar	8
Change Administrator Password Screen	15
The Database Server Name	23
Using A Different External Host Name For LISTSERV	27
Global Settings for the Maestro User Interface	29
General Administration of the Maestro User Interface - General Settings	30
General Administration of the Maestro User Interface - Advanced Security Options	30
General Administration of the Maestro User Interface - Runtime Administration	31
Database Plugins	31
System Database Connection	32
User Rights	33
LISTSERV Web Interface Access	33
Default URL Settings	34
Default LISTSERV Connection	37
Default LISTSERV Connection - Hosted LISTSERV Lists Settings	38
Default LISTSERV Connection - Automatic Bounce Handling	39
Default Size Limits	39
Default Content Restrictions	40
Default Recipients Restrictions - Recipients Type Restrictions	41
Default Recipients Restrictions - Recipients Upload Restrictions	42
Default Tracking Restrictions	42
Default Hosted Data Settings - General Settings	43
Default Hosted Data Settings - Hosted List Settings	43
Default Hosted Data Settings - Hosted Data Import Restrictions	44
Default DomainKeys Settings	45
Default Auto-Archive Settings	46
Automatic Login Link	49
Administrative Email Notifications	56
Administering User Accounts and Identities	59
Defining User Account	60
Defining Identity	61
User Rights Management Screen	62
Editing User Account Settings	63
User Account and Group Menus	64
User Right Settings	65
URL Settings	67
Size Limits	69
Job ID Prefix Settings	70

Content Restrictions	70
Drop-In Content Restrictions	71
Recipients Restrictions - Type Restrictions	72
Recipients Restrictions - Upload Restrictions	72
Tracking Restrictions	73
DomainKeys Settings	75
Auto-Archive Settings	76
External Access Settings	76
Deleting an Account	78
Deleting an Identity	78
Deleting a Group	79
The Toolbar in the LISTSERV Maestro User Interface	81
The Dashboard	91
Refreshing the Dashboard	92
The Edit the Dashboard Name Screen	92
The Hide/Show Default Section Screen	93
System Metrics on the Dashboard	94
The System Metrics Report - Request Counts Tab	95
Downloading the System Metrics Report	96
Renaming the Dashboard Report	97
Copying a Dashboard	98
Creating a Dashboard	99
Switching Dashboards	100
Archiving a Job from the Completed Jobs Screen	102
Importing a Job from the Archived Jobs Screen	103
Change Job Owner from Job Owner Tab	105
Change Report Owner Screen	105
Example: Changing the Owner of a Sender Profile	108
Change Target Group Owner Screen	111
Change Ownership of a Target Group Category	111
Change Ownership of an Individual Target Group	112
Change Ownership of a Recipient Dataset	113
Recipient Datasets Administration Screen	113
The Hosted List Administration Screen	114
JDBC Driver Layers	121
ODBC Plugin Layers	121
Adding a Database Plugin	122
System Database Connection Screen	125
System Database Connection Details Screen for DB2	125
The General Component Settings for Administration Hub Screen	132
Component Communication Pathways	160
Multiple Logins	165
Example of Base64 Encoded Outfile	176
Imported Certificate	178
Example of Recipients Profile Data Table	186
Example of Recipients ID in Data Table	187

The DomainKeys Settings Screen	202
LISTSERV Web Interface Links	206
Editing an Existing Link	207
LISTSERV Web Interface Mappings	210
Creating a New Account Mapping	211
Sample Proxy Setup	233


List of Tables

- Restrictions for Hosted LISTSERV List Support 19
- Backup History 134
- Supported SSL Ciphers 183
- Maestro User Interface INI-File Entries 190
- Administration Hub INI-File Entries 195
- Maestro Tracker INI-File Entries 196
- Basic Configuration for Tomcat INI-File Entries 197
- Advanced Configuration for Tomcat INI-File Entries 199
- Supported Charsets 253

Preface - About This Manual

Every effort has been made to ensure that this document is an accurate representation of the functionality of LISTSERV® Maestro. As with every software application, development continues after the documentation has gone to press so small inconsistencies may occur. We would appreciate any feedback on this manual. Send comments via email to: MANUALS@LSOFT.COM

The following documentation conventions have been used in this manual:

- Menus, options, icons, fields, and text boxes on the screen will be bold (e.g. the **Administer User Accounts** icon).
- Clickable buttons will be bold and within brackets (e.g. the **[OK]** button).
- Clickable links will be bold and underlined (e.g. the **Edit** link).
- Directory names, commands, and examples of editing program files will appear in `Courier New` font.
- Emphasized words or phrases will be underlined.
- Some screen captures have been cropped for emphasis or descriptive purposes.
- Unless otherwise specified, directory paths are for Microsoft Windows installations of LISTSERV Maestro. For Linux or Solaris, substitute the Maestro top-level directory for `\Program Files\L-Soft\Application Server`.
-  This symbol denotes a new feature for LISTSERV Maestro 4.0.

What's New in LISTSERV[®] Maestro 4.0

This section highlights the new features for the **System Administrator** in LISTSERV[®] Maestro 4.0.

Switching Between Interfaces

A new single sign-on feature is now available for the administrator, allowing one-click access to both the HUB and LUI. For details, see Section 1.4 [Accessing the Hub Administrator Interface](#).

New Look and Feel

The Administration HUB now has a new look and feel that makes it easy to perform administration tasks. For details, see Section 1.5 [Understanding the Hub Administrator Interface](#).

Enabling Action Tracking

You can now track actions that a user performs on a “target” website. For details on enabling for the User Interface, see Section 5.2 [Application Default Settings](#) and Section 7.4.2 [Editing Component Specific Settings for Single and Group Users](#).

Enabling the Job Data Export Feature

It is now possible to export a specific job or range of jobs depending on access parameters. The export is triggered by making a request to an external access URL and results in an export file in XML format. For details on enabling the external access setting for a user account or group, see Section 7.4.2 [Editing Component Specific Settings for Single and Group Users](#).

Monitoring Your System Metrics

LISTSERV Maestro now monitors your system metrics, which gives you access to status reports and helps you monitor the amount of free disk space on the installation disk. For details on enabling, see Section 5.1 [Application Settings](#). For details on using this new feature, see Section 8.2.2 [System Metrics](#) and Section 8.2.3 [System Metrics Reports](#).

Defining the Hosted Data Import Restrictions

The data administrator can now save predefined “recipient importers” and import recipients directly into the Recipient Warehouse. Because of this, the Hosted Data settings have been updated. For details, see Section 5.2 [Application Default Settings](#) and Section 7.4.2 [Editing Component Specific Settings for Single and Group Users](#).

Behavioral Changes to the Toolbar

In the User Interface, the behavior of the **Dataset** menu on the Toolbar has been updated to accommodate those who may need to access this menu while working with lists. Because of this, there are several changes to the Toolbar. For details, see Section 8.1.1 [Behavioral Changes to the Toolbar](#).

SSL Cipher Support

SSL ciphers are now supported by LISTSERV Maestro. For details, see Section 18.3.9 [Supported SSL Cipher Suites](#).

Mixing SSL and Non-SSL Access on One Server (Mixed Mode)

It is now possible to secure some users/groups with SSL (using https://) while other users/groups are not secured (using http://). To facilitate this new feature, several changes had to be made to the previous admin settings and configurations. Because of this, please review this manual in its entirety if you are configuring a Maestro instance. Some main topics you should pay close attention to are:

- Section 14.2.1 [Configuring the HTTP Port](#)
- Section 14.2.2 [Configuring the HTTPS Port](#)
- Section 20 [Editing LISTSERV Maestro INI Files](#)
- Section 18 [Securing Access with SSL](#), with emphasis on
 - Section 18.3.6 [Making LISTSERV Maestro Aware of the Server Certificate](#) ,
 - Section 18.3.7 [Securing a Server with Multiple Host Names](#) ,
 - and the new Section 18.3.8 [Mixing SSL and Non-SSL Access on One Server \(Mixed Mode\)](#)

URL Settings and Other Hosted Data Changes

The Tracking URL screen has been replaced by the URL Settings screen, which is where you define the various URLs used by the User Interface. Because of this new screen, please note the following changes as well:

- To define the subscriber pages URL, instead of using the Hosted Data Settings screen, you will now use the new URL Settings screen.
- Previously, when you wanted to define the login restrictions host name you would use the Login Restrictions screen. This screen has been removed, and any login restrictions are now defined on the new URL Settings screen.

For details on this new screen, see Section 5.2.1 [Default URL Settings](#) and Section 7.4.2 [Editing Component Specific Settings for Single and Group Users](#).

Emergency Admin HUB Access

Under certain circumstances, the administrator may not be able to login to the Administration HUB using the normal Maestro User Interface login page. If this happens, and the problem stems from an incorrect configuration in the HUB, then the administrator can now have direct login access to the HUB in order to fix this problem. For details, see Section 31 [Emergency Admin HUB Access](#).

Moving a LISTSERV Maestro Installation

Sometimes it is necessary to replace the hardware on which a LISTSERV Maestro instance is installed, or to move a LISTSERV Maestro installation from one server to another. For details, see Section 23 [Moving a LISTSERV Maestro Installation](#).

Configuring Maestro to Bind to Different HTTP Ports on Different IP Addresses

You can now configure LISTSERV Maestro to bind to different ports on different IP addresses, if necessary. For details, see Section 15 [Defining the IP Addresses and HTTP Ports](#).

Sharing a Server with IIS

You can now configure LISTSERV Maestro to act as a backend server behind IIS. For details, see Section 32 [Sharing a Server with IIS](#).

Configuring Aliases for Access URLs

You can now configure aliases for the access URLs in the Administrative HUB. For details, see Section 5.2.1 [Default URL Settings](#) and Section 7.4.2 [Editing Component Specific Settings for Single and Group Users](#).

Section 1 Introduction to Maestro Administration

Designed specifically to work with LISTSERV® 15.5, LISTSERV Maestro 4.0 allows users to easily create and send personalized email messages using a web interface. Incorporated into this powerful tool is a tracking component that can collect data every time a recipient opens an email message or clicks on a URL contained within the message.

The LISTSERV Maestro program is comprised of three components that work together:

- **The Administration Hub** – Controls all user and program settings. It is the central component that stores registry and account information. It is accessed both by the Maestro User Interface and by Maestro Tracker to validate login information. It has its own administrator user interface.
- **The Maestro User Interface** – The actual user interface. Individuals and groups use it to create and distribute customized email messages. It is also used to access, view, and download the collected tracking data, and to maintain the recipient repository (datasets and lists).
- **The Maestro Tracker** – Receives and compiles tracking data from delivered email messages.

In addition to LISTSERV Maestro's three components, LISTSERV Maestro 4.0 also relies on the existence of an installation of **LISTSERV 15.6**. LISTSERV receives email jobs from LISTSERV Maestro and prepares them for delivery. It is also used to process bounced mail. In certain cases, LISTSERV may also act as an interface between LISTSERV Maestro and an external DBMS. In addition, if both applications are configured correctly, you can switch between the two with a single sign-on, meaning you won't have to logout of one and then login to the other.

By default, LISTSERV Maestro 4.0 comes installed with MySQL® 4.1.7 as the internal system database. Use of this internal system database is optional. It is possible to decide during installation not to install the internal MySQL database, in which case some external database must be installed instead. It is also possible to install the internal database during the initial setup, and then switch to some external database later.

Please see <http://www.mysql.com> for details about MySQL 4.1.7 and its features to determine if it meets your organization's requirements and expectations for a production database. Reviewing the features will help to decide if the internal database should be used or if a different external system database should be installed.

LISTSERV Maestro can use the external database to store its own system data as well as to select recipient lists from database tables and drop-in content elements. LISTSERV Maestro can connect to several databases in this way. Supported databases are:

- Oracle® 8i, Oracle® 9i, Oracle® 10g, and compatible versions
- DB2® V7.2, V8.2, and compatible versions
- MySQL® 4.x and compatible versions, as well as 3.23.42 and later 3.23x builds
- SQL® Server 7.0 and 2000

LISTSERV and LISTSERV Maestro also require access to one or more SMTP servers to perform the actual delivery of email jobs for LISTSERV Maestro and LISTSERV. Any standards-compliant SMTP server will work.

The three LISTSERV Maestro components, LISTSERV, the SMTP server(s), and the optional external database may be installed on any combination of hosts, from one single host shared by all components to six or more dedicated hosts, one for each component (it is possible to have multiple LISTSERV servers and multiple SMTP servers). If different components are installed on separate servers, it is not necessary that all of the servers have the same operating system. It is possible to install the Maestro User Interface and Administration Hub components on a Windows server and at the same time the Maestro Tracker component on a Linux server (or other combinations). For more information on host restrictions, installing LISTSERV Maestro, and starting and stopping the LISTSERV Maestro service, see the LISTSERV Maestro Installation Manual.

1.1 Maestro User Roles

It is important to understand the different roles involved in administering and using the various components of a LISTSERV Maestro system. In a small organization, the same person may play many of these roles. In larger organizations, the following duties will likely be distributed among several different people:

- **System Administrator** – Responsible for the installation and initial configuration of the LISTSERV Maestro applications. The system administrator must have Administrator or root access to the computer(s) on which LISTSERV Maestro applications will reside.
- **LISTSERV Maestro Administrator** – Responsible for administration of LISTSERV Maestro through the HUB component and the “admin” account in the Maestro User Interface. The LISTSERV Maestro Administrator acts as the “master account” for all LISTSERV Maestro users.
- **Data Warehouse Administrator** – Responsible for administering recipient data within LISTSERV Maestro. See the “LISTSERV Maestro Data Administrator’s Manual” for details.
- **LISTSERV Site Administrator** – Responsible for the configuration and administration of LISTSERV (including configuring LISTSERV for database access, adding LISTSERV “postmaster” accounts, and so on). See the “LISTSERV Site Manager’s Manual” for details.
- **Database Administrator** – Responsible for the initial installation of the external database(s). Also responsible for monitoring available space in the database(s), database performance tuning, routine database backups, and other routine database maintenance tasks.
- **Maestro User** – Responsible for creating, sending and tracking email jobs through LISTSERV Maestro. There are many different user responsibilities, which may fall to different users. See the “LISTSERV Maestro User Guide” for details.
- **SMTP Server Administrator** – Responsible for the administration of the SMTP email delivery engine.

For many organizations, it may be desirable to have some overlap among the various roles. For instance, the Data Warehouse Administrator may also be a regular Maestro User, or the System Administrator and LISTSERV Administrator may be the same person. In other cases, there may be clear distinctions between some of the responsibilities (e.g., it may not be desirable for a Maestro User to also have System Administrator access to the computer running LISTSERV Maestro).



Tip: It is usually a good idea to understand and assign the various user roles to groups and individuals within your organization before even installing LISTSERV Maestro. This may save considerable time and confusion.

1.2 Remote Administration Access

The administrator can access log files or query the current version remotely.

1.2.1 Remote Log Access

The three main LISTSERV Maestro components all write their own log files. These files are in the “logs” subfolder of each component’s home folder inside the installation folder. However, in some situations the administrator does not have access to these folders, but still wants to access the log files. To solve this, LISTSERV Maestro offers remote log file access. The remote access allows an administrator to download the log files from the server, directly in the web-browser.

Before you can access the log files of a component, you first have to configure the component for remote log access. To do so, edit the INI-file of the component and add the following entry:

```
RemoteAdminPassword=PASSWORD
```

where you replace “PASSWORD” with a password only known to authorized administrators (although, for security reasons, you should not use your normal admin password from the Administration Hub).



Note: Since the password will later be used as a parameter in a URL, you should only use URL-safe characters in the password (e.g. you are on the safe side if you only use alphanumeric characters).

Remember that you have to add this entry to each component’s INI file, e.g. to lui.ini, hub.ini, and tracker.ini. If you do not add the entry to one of the INI files, then you will not be able to access the log files of that component (but you will still be able to access logs of the other components where you have added the entry).

To disable remote log access, remove the entry from the INI file(s) or comment it out.

Whenever you add, change, remove or out-comment this entry, the change will be effective immediately, e.g. you do not have to restart the component!

Once you have configured the component(s) for remote log access, you can access their log files from any web-browser on any computer that has HTTP access to the component in question. You only need to know the “PASSWORD” you configured in the INI file(s) and the day of the log file you want to access (for normal log files) or the backup-ID (for backup log files):

- To view a Maestro User Interface log file, access the following URL:

```
http://HOST:PORT/lui/downloadLog
```

- To view a Administration Hub log file, access the following URL:

`http://HOST:PORT/hub/downloadLog`

- To view a Maestro Tracker log file, access the following URL:

`http://HOST:PORT/trk/downloadLog`

- To view a Backup log file, access the following URL:

`http://HOST:PORT/hub/downloadLog`

where you replace:

“HOST” with the host name of the server running the component you want to access.

“PORT” with the HTTP-port on that server (“PORT” can be left out if the HTTP-port is “80”),

On the page that is displayed, enter the remote admin password that you have configured in the INI-file (see above) into the **Password** field and click **[Apply]**.

If you enter only the password, then the displayed log file will be the one of the current day. If you want to see a log file of a different day, enter the date of the day you want to view into the first **Date** field (leave the second **Date** field empty). If you want to see the log files of a range of days all at once, enter the date of the first day in the first **Date** field and the date of the last day in the second **Date** field. (For all date input, use the format “YYYYMMDD”, where “YYYY” is the year with 4 digits, “MM” is the month with 2 digits and “DD” is the day of the month with 2 digits.) Each time you change a date setting, click **[Apply]** to make it effective.

You can also enter a text into the **Search String** field (and click **[Apply]**). This will have the effect, that all occurrences of this string in the log file will be highlighted and all log entries which do not contain the search string will be initially hidden so that you can concentrate on the log entries that contain the search string. The hidden areas are marked with a little “+” symbol at the left margin. Click the “+” to unfold a hidden part and make it visible again. The left margin will then show a ruler with “-“ symbols at the top and the bottom of the unfolded area. Click on the ruler or the “-“ symbols to hide the area again.



Note: If you need to provide the log file to a 3rd party (for example, to L-Soft support), then you can use the **Save as plain text** link at the bottom right of the page to download and save the backup file in its plain text format. Please supply this plain text log file to L-Soft support if necessary (i.e. please do not simply do a copy & paste from the browser and do also not a “Save As...” of the log-viewing page, as this would also mean that all the additional HTML data used for formatting the log-view would be included).

If you access the log file of the Administration Hub component, then the page also has an additional **Backup ID** field. You can use this field to access one of the backup logs of the HUB: Enter the backup-ID of the backup for which you want to view the log and make sure that the two **Date** fields are empty (as long as the **Date** fields contain a value, the **Backup ID** is ignored and the standard log is retrieved). Then click **[Apply]**.

Initially, the page displays the backup log with all the backup details hidden and only the general information visible. As explained above, the hidden parts are marked with a little “+” symbol at the left margin and can be unfolded for viewing. Once unfolded, they can also be hidden again.

The backup-ID is a sequence of digits and letters and is unique for each backup. You can find the information about which ID a given backup has by looking at the normal HUB log file: At the backup start-time you will find an entry like “BackupMaster starts backup (Backup ID: xxxxxx)” where “xxxxxx” is the backup-ID that you need to enter in the above **Backup ID** field.

1.2.2 Remote Version Query

The current version and build number of all components can be queried remotely. This is done with a simple HTTP-request - a URL typed into the address field of any browser. The result of the query will be displayed in the browser. This query can also be used to verify that a fresh installation is operational.

- To query LUI: `http://LUISEVER/lui/build`
- To query HUB: `http://HUBSERVER/hub/build`
- To query TRK: `http://TRKSERVER/trk/build`
- To query LUI's list context: `http://LUISEVER/list/build`

Substitute the name of the Maestro User Interface server, Administrative Hub server, or Maestro Tracker server for LUISEVER, HUBSERVER, and TRKSERVER (respectively).

1.3 Client System Requirements

Depending on the operating system of the client used for access, the following browsers are supported when accessing the Maestro User Interface or Administration Hub:

- Client with Windows – Microsoft® Internet Explorer 5.5 or later, Mozilla® 1.0.0 or later, other browsers based on a compatible Mozilla version (e.g., Firefox 1.0 or later, Netscape® 7.0 or later).
- Client with Linux – Mozilla® 1.0.0 or later, other browsers based on a compatible Mozilla version (e.g., Firefox 1.0 or later, Netscape® 7.0 or later).
- Client with Mac OS X – Mozilla Firefox or the Mac native Safari browser (version 2.0, build 412 or later).

To access the Maestro User Interface or the Administration Hub, we strongly recommend that only Windows, Linux, or Mac OS X is used with the browsers and browser versions listed. Other operating systems, browsers, or browser versions are *not* supported.

The client does not necessarily have to have the same operating system as the LISTSERV Maestro server. A Linux client can be used to access LISTSERV Maestro on a Windows server and vice versa. Similarly, the different components of LISTSERV Maestro may run on different operating systems, if they are installed on separate servers. For example, the Maestro User Interface and Administration Hub components may be installed on a Windows server, while the Maestro Tracker is installed on a Linux server.



Important: Recipients of email being tracked by LISTSERV Maestro may use whatever browser they wish to access the URLs contained in the messages sent by LISTSERV Maestro. Tracking occurs no matter which browser is used by email recipients.

1.4 Accessing the Hub Administrator Interface



With the new Single Sign-On feature, the LISTSERV Maestro Administrator can now switch back and forth between the Administration HUB and User Interface with one click. No more hassle, no more logging in and out, just one simple click is all it takes.

To access LISTSERV Maestro, open a web browser and access the URL:

```
http://<YOUR_MAESTRO_USER_INTERFACE_HOST>/lui
```

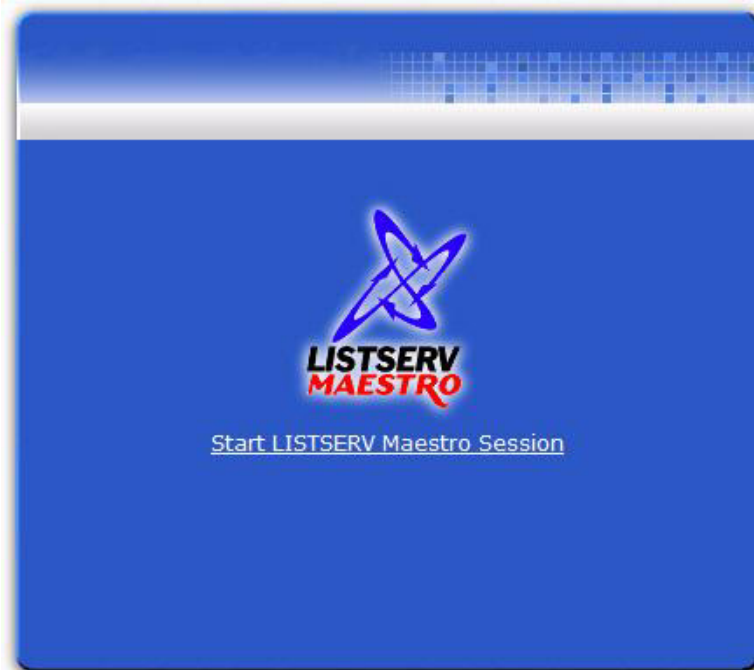
In the URL above, you need to replace “<YOUR_USER_INTERFACE_HOST>” with the hostname of the server running the Maestro User Interface component (LUI). In addition, if the server uses a different HTTP-port than the standard port “80”, then you need to include this different port with the host name, separated with a colon “:”, like this: “http://HOSTNAME:PORT/lui”.



Tip: In a Windows installation, a shortcut for this address will appear in the Windows Start Menu under **Programs > L-Soft Application Server**.

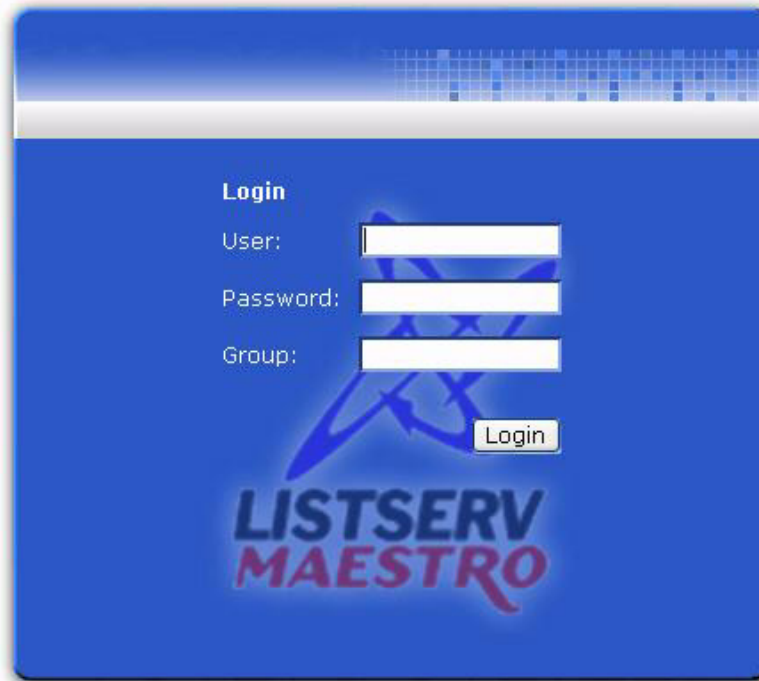
This URL provides access both to the Maestro User Interface (for normal users and the administrator) and to the Administration Interface (for the administrator only); in addition, the URL redirects you to the new login area.

Figure 1-1 Accessing LISTSERV Maestro



From here, click on the **Start LISTSERV Maestro Session** link to open the Login Screen.

Figure 1-2 LISTSERV Maestro Login Screen



Enter the Administrator's user name and password (the user name is *admin* and the password is established during installation). Click the **[Login]** button to enter the Administration HUB.



Note: The Administrator's account does not belong to a group; therefore, the **Group** field should be left empty.




1.4.1 Switching Between Interfaces

With the new "single sign-on" feature, the LISTSERV Maestro Administrator can now switch back and forth between the Administration HUB and User Interface with one-click. No more hassle, no more logging in and out, just one simple click is all it takes.

To switch to the LISTSERV Maestro User Interface, simply click the **Switch to LUI** link at the top of the HUB screen.

Figure 1-3 Switching Between Interfaces



Note: If you switch to the User Interface and then logout from there, the next time you (as the administrator) login, you will taken directly to the User Interface. Similarly, if you logout of the HUB, the next time you login, you will be taken directly to the HUB.

1.5 Understanding the Hub Administrator Interface

Administering LISTSERV Maestro involves many different tasks and interaction with more than just the Administration Hub (HUB). Administrators will have to understand how LISTSERV Maestro is situated within the institution's infrastructure. This understanding is critical for making decisions about settings for all the application's components to ensure consistency and compatibility with new or existing systems. Consequently, this manual is organized around those different tasks an administrator needs to perform in order to set up, monitor, backup, and change an installation of LISTSERV Maestro. It also serves as a reference for advanced systems configuration, touching on the HUB Interface screens as they fit in to each task.



The Accounts and Identities screen is the opening screen of LISTSERV Maestro's Administrator Interface. This screen gives you quick and immediate access to one of the most important administrative duties – managing user accounts.



1.5.1 The Toolbar

The Toolbar contains menus and icons that give you quick access to the different functions in LISTSERV Maestro, letting you easily navigate from one area to another. It appears on each LISTSERV Maestro page and displays the currently available options.



Note: There are quite a few different sub-menus and options that may appear in the various parts of the main menu. However, not all of these options and sub-menus are always visible. The main menu and its sub-menus only show those options and menus that are currently available according to the context of the page you are on.

Figure 1-4 The Toolbar



The **Accounts and Identities** menu contains options for the creation and administration of user accounts and identities. The possible options are:

- **Accounts and Identities Overview** – Display the Accounts and Identities screen.
- **New Account** – Create a new user account.
- **New Identity** – Create a new identity.

The **User Account** menu is only available when a user is selected. It contains options to administer the currently selected user account. The possible options are:

- **Edit** – Edit the account name, group, identity, and/or password, and define if the account is allowed to change its password or not.
- **Delete** – Delete the user account.
- **User Right Settings** – Define the user rights for the selected account.
- **LISTSERV Connection** – Define the LISTSERV connection settings for the selected account.
- In addition, if the selected account is a single (non-group) user account, then all of these options are also available:

- **URL Settings** – Define the settings for various URLs used by LUI for the selected account.
- **Size Limits** – Define the message size limit for the selected account.
- **Job ID Prefix** – Define the job-ID prefix for the selected account.
- **Lite Mode Restrictions** – Define the lite mode restrictions for the selected account.
- **Content Restrictions** – Define the content restrictions for the selected account.
- **Recipients Restrictions** – Define the recipients restrictions for the selected account.
- **Tracking Restrictions** – Define the tracking restrictions for the selected account.
- **Hosted Data Settings** – Define the hosted data settings for the selected account.
- **DomainKeys Settings** – Define the DomainKeys settings for the selected account.
- **Auto-Archive Settings** – Define the auto-archive settings for the selected account.

The **Group** menu is only available when a group is selected. It contains options to administer the currently selected group. The possible options are:

- **Delete** – Delete the selected group and all accounts in it.
- **LISTSERV Connection** – Define the LISTSERV connection settings for the selected group.
- **URL Settings** – Define the settings for various URLs used by LUI for the selected group.
- **Size Limits** – Define the message size limit for the selected group.
- **Job ID Prefix** – Define the job-ID prefix for the selected group.
- **Lite Mode Restrictions** – Define the lite mode restrictions for the selected group.
- **Content Restrictions** – Define the content restrictions for the selected group.
- **Recipients Restrictions** – Define the recipients restrictions for the selected group.
- **Tracking Restrictions** – Define the tracking restrictions for the selected group.
- **Hosted Data Settings** – Define the hosted data settings for the selected group.
- **DomainKeys Settings** – Define the DomainKeys settings for the selected group.
- **Auto-Archive Settings** – Define the auto-archive settings for the selected group.

The **Global Settings** menu contains options to administer the global configuration settings. The possible options are:

- **Administration Hub** – Define the global configuration settings for the Administration Hub component. The sub-menu may contain the following options:
 - **General Administration** – Configure general Administration Hub related settings.
 - **Administrative E-mail Notifications** – Configure the email notification settings (for all components).
- **Maestro User Interface** – Define the global configuration settings for the Maestro User Interface component. The sub-menu may contain the following options:
 - **General Administration** – Configure general Maestro User Interface related settings.
 - **Database Plugins** – Configure the available database plugins.
 - **System Database Connection** – Configure the system database connection.
 - **User Rights** – Configure the Maestro User Interface specific user rights of all user accounts.
 - **LISTSERV Web Interface Access** – Define the configuration for direct access to the LISYSERV Web Interface (WA) from the LISYSERV Maestro Interface (and vice versa). The sub-menu may contain the following options:
 - **LISTSERV Web Interface Links** – Configure links between the LISYSERV Maestro User Interface and one or several LISYSERV Web Interfaces (WAs).
 - **LISTSERV Web Interface Account Mappings** – Configure mappings between LISYSERV Maestro accounts and LISYSERV Web Interface (WA) accounts, to allow a user to move between the two interfaces without having to login again.
 - **Default URL Settings** – Define the default settings for various URLs used by LUI.
 - **Default LISYSERV Connection** – Define the default LISYSERV connection settings.
 - **Default Tracking URL** – Define the default tracking URL.
 - **Default Size Limits** – Define the default message size and file upload size limits.
 - **Default Lite Mode Restrictions** – Define the default lite mode restrictions.
 - **Default Content Restrictions** – Define the default content restrictions.
 - **Default Recipients Restrictions** – Define the default recipients restrictions.
 - **Default Tracking Restrictions** – Define the default tracking restrictions.
 - **Default Hosted Data Settings** – Define the default hosted data settings.

- **Default DomainKeys Settings** – Define the default DomainKeys settings.
- **Default Auto-Archive Settings** – Define the default auto-archive settings.
- **Maestro Tracker** –Sub-menu for global configuration settings for the Maestro Tracker component. The sub-menu may contain the following options:
 - **General Administration** – Configure general Maestro Tracker related settings.

The **Utility** menu contains general administer options. The possible options are:

- **Change Password** – Change the password of the administrator account.
- **Print Page** – Print the current page/screen.
- **About LISTSERV Maestro** – Show LISTSERV Maestro version information.

The **Logout** option allows you to log out from LISTSERV Maestro (after first confirming that you really want to log out).



The **Help** icon is used to access the help associated with the current screen.

Section 2 Configuring LISTSERV Maestro for First Use

After installation, it is necessary to execute a few initial configuration steps in the Administration Hub before LISTSERV Maestro can be fully used. Access the Administration Hub as described in Section 1.4 [Accessing the Hub Administrator Interface](#) and log in with the administrator password.

1. **Configure the System Database** (Optional) – The Maestro User Interface component of LISTSERV Maestro uses a “system database” to store its working data – recipient profiles, job ID numbers, tracking information, and so on.

An “internal” database (based on MySQL) is included as part of the application and may be used as the system database. Using this internal database will allow the application to run “out-of-the-box”. An optional external database may be configured in place of the default internal database if desired. Switching the system database from internal to external and vice versa can be done at a later time, if necessary. For more information on configuring an external database, see Section 10 [The System Database](#).

2. **Define User Database Connections** (Optional) – The Maestro User Interface may optionally access “user databases” to retrieve information to build recipients lists in the Recipient wizard or Target Groups wizard, or to create drop-in content elements.

If a user database is going to be used, LISTSERV Maestro must be configured to access it. The appropriate driver must be installed on the server running the Maestro User Interface (the LUI component) and the appropriate “plugin” must be registered in the Administration Hub (the HUB component). For more information on configuring an external database, see Section 9 [Defining External Database Connections](#).

3. **Set up LISTSERV** – See Section 4 [LISTSERV and LISTSERV Maestro](#) for details on how to set up LISTSERV to work with LISTSERV Maestro.

On Windows, if LISTSERV Maestro was installed using the **Express Setup** option from the Setup Suite Installation Kit, then this step can be skipped, unless you want to make additional configurations (for example, to enable Hosted LISTSERV Lists).

4. **Configure the Maestro User Interface (LUI)** – Two steps in particular need to be accomplished. (These are also handled automatically by the **Express Setup** option on Windows. If you used this option during installation, then you may skip this step, unless you want to make changes to the default setup.)

- **Define the default LISTSERV Connection** – This connection is used for all accounts that do not have individual connection parameters configured. If a single LISTSERV connection is shared among all users, then configure this connection as the default connection and leave the configuration parameters of individual users (or groups) empty. Leave the default connection parameters empty only if connection parameters for all users and groups on the account or group level will be configured individually. See Section 5 [Settings for the](#)

[Maestro User Interface](#) for more information on configuring the default LISTSERV connection.

- **Define the tracking URLs** – If tracking is to be used, it is necessary to define the domain name of the tracking server that will be used in the URLs for tracked links. Leave the default tracking URL parameter empty only if this parameter will be defined individually for all groups. See Section 5.3 [Setting the Default Tracking URL](#) for information on how to define the default tracking URLs.
5. **Configure Global Component Settings** – Establish administrative policies and procedures and configure global component settings to reflect these. In particular decide upon:
 - **Backup procedures** – Decide what time to make daily backups, how many backups to keep, where the backup(s) will be saved, whether any external processes will run before or after a backup.
 - **Archival procedures** – Decide the circumstances under which old jobs will be archived. Define the archive folder.
 - **Taking the system down** – When will the system be taken down for maintenance, how will users be warned, what restrictions will be imposed before the system is taken down.
 - **User account restrictions** – What, if any, restrictions will be imposed on user accounts based upon the type(s) of user, size limits of uploads, recipients types, drop-in content types, and so on.
 - **Tracker event transfer frequency** – Decide how often the tracking information needs to be refreshed. This depends on how current the tracking reports need to be. The default time period is 10 minutes.
 6. **Create User Accounts and Identities** – Create at least one user account that can be used to access the Maestro User Interface. See Section 7 [Creating and Editing User Accounts and Identities](#) for more details on how to proceed with creating user accounts, changing the password for the “sample” user account, or deleting the “sample” account.
 7. **Secure the Default Keystore** – If any of the LISTSERV Maestro components are going to be secured using SSL (Secure Socket Layer), then change the password for the default keystore for trusted root certificates that is shipped with Java. The instructions for this procedure are located in Section 18.3.3 [Securing the Trusted Root Certificate Keystore](#).

Section 3 Changing the Administrator Password

To change the administrator password, click on the **Utility** menu and select the **Change Password** option. The Change Administrator Password screen opens.

Figure 3-1 Change Administrator Password Screen



The screenshot shows a dialog box titled "Change Administrator Password" with a lock icon. It contains three text input fields labeled "Old Password:", "New Password:", and "Confirm Password:". At the bottom, there are two buttons: "OK" and "Cancel".

Enter the old password in the top field and a new password in the field beneath. Confirm the new password by retyping it in the third field, and then click the **[OK]** button to record the changes or the **[Cancel]** button to disregard.

Section 4 LISTSERV and LISTSERV Maestro

LISTSERV Maestro uses an instance of LISTSERV to send email jobs that are prepared in the Maestro User Interface. Any instance of LISTSERV, or multiple instances of LISTSERV, can be used to send jobs as long as each instance has a LISTSERV Maestro license, is reachable over the network, has been configured to accept jobs from LISTSERV Maestro, and LISTSERV Maestro has been configured to connect to it. A hierarchy of LISTSERV Connection settings can be used to configure a global application default connection, a default connection for each group, a separate sender address for group users, and a default connection for each single user not in a group. For more information on configuring individual and group user settings, see Section 7.4.2 [Editing Component Specific Settings for Single and Group Users](#).

The system works by having the Maestro User Interface send a “distribute job” to LISTSERV. A DISTRIBUTE job is a series of commands to LISTSERV that essentially says, “Take this message and send it to these recipients.” In order to successfully process a DISTRIBUTE job, LISTSERV needs to have a list of recipients (email addresses), and the message itself. A complete DISTRIBUTE job must include one or more command lines giving instructions to LISTSERV and an authenticating password. LISTSERV Maestro handles these and many other steps automatically.

Normally, LISTSERV does not accept “distribute” commands from everyone. LISTSERV accepts this command only if the “distribute job” is sent from an account (email address) that is configured in LISTSERV to have the right to DISTRIBUTE jobs. The reason for this is to avoid allowing LISTSERV to be hijacked for spamming and other unethical purposes.

LISTSERV Maestro may also access an instance of LISTSERV to create and manage Hosted LISTSERV lists. This feature utilizes nearly all of the features of traditional LISTSERV lists in combination with the LISTSERV Maestro user interface. It also requires additional configuration within LISTSERV to allow LISTSERV to connect to the user database and to allow LISTSERV Maestro to create new LISTSERV lists.

4.1 Preparing LISTSERV for LISTSERV Maestro

Several steps are required in order to prepare LISTSERV for communication with LISTSERV Maestro. These include preparing LISTSERV to accept DISTRIBUTE jobs from Maestro, preparing LISTSERV to allow LISTSERV Maestro to create new LISTSERV lists, preparing the LISTSERV list archive folder, and preparing LISTSERV for database access. These steps require access to the LISTSERV Site Configuration file on the LISTSERV server, and need to be carried out by the LISTSERV Site Administrator. For more information on LISTSERV’s Site Configuration file, see the LISTSERV Site Manager’s Manual at <http://www.lsoft.com/resources/manuals.asp>.

4.1.1 Preparing LISTSERV to Process DISTRIBUTE Jobs from LISTSERV Maestro

In order to process email jobs from LISTSERV Maestro, LISTSERV needs to be prepared in certain ways. Authentication between LISTSERV and LISTSERV Maestro

happens by way of an email address and password combination. This means that LISTSERV needs to have an email address and password stored for LISTSERV Maestro in order to give LISTSERV Maestro the right to send DISTRIBUTE jobs. (This email address and password combination will later be entered in to the Administration Hub; see Section 5 [Settings for the Maestro User Interface](#).)

For Windows, if the **Express Setup** option was used to install the full LISTSERV Maestro Suite, then this step is handled automatically for the default LISTSERV connection. You can skip this step unless you want to add more distribute accounts (i.e. so that each group uses a different account, which is recommended if you do not want separate groups sharing data).

Prepare LISTSERV for processing email jobs from LISTSERV Maestro as follows:

- In LISTSERV, configure an email address to have the DISTRIBUTE right. The email address must be defined in LISTSERV's Site Configuration file with the following parameter:

`DIST_ALLOWED_USERS`: This parameter confers only the right to send DISTRIBUTE jobs and is the recommended one to use. An example from a Windows `SITE.CFG` file might look as follows:

```
DIST_ALLOWED_USERS=MAESTRODIST@EXAMPLE.ORG
```

Remember that LISTSERV must be restarted in order for any site configuration changes to take effect.

- In LISTSERV, register a password for the email address that was specified under the `DIST_ALLOWED_USERS` or `POSTMASTER` setting. There are two ways to do this:
 - The first way is to send an email message to `LISTSERV@EXAMPLE.ORG` (substituting the domain name of the LISTSERV server in place of `EXAMPLE.ORG`) from the newly registered email address with the following command in the message body:

```
PW REP newpassword
```

where `newpassword` represents the desired password. LISTSERV will send a confirmation message back by email. You must reply to that confirmation message with another message saying `OK`. If the LISTSERV Web Interface is running on the LISTSERV server, then the request can be confirmed by visiting the URL provided in the confirmation message.

- The second way to do this is to send an email message to `LISTSERV@EXAMPLE.ORG` (substituting the domain name of the LISTSERV server in place of `EXAMPLE.ORG`) from the LISTSERV site administrator's address (an address defined as a `POSTMASTER` in the LISTSERV Site Configuration file) with the following command in the message body:

```
PWC ADD email newpassword PW=createpw
```

where `email` is the email address being used for the LISTSERV Maestro jobs,

`newpassword` is the password being assigned to that email address, and `createpw` is the password defined in the Site Configuration file as the `CREATEPW` or the personal password registered for the postmaster address.

By configuring LISTSERV with these steps, an email address and password combination now exists that has the right to send DISTRIBUTE jobs. This address and associated password are necessary for LISTSERV to accept a DISTRIBUTE job from LISTSERV Maestro. The configured address, referred to as the "sender" address, is the sender of the DISTRIBUTE job. It is not the same email address that is used in the Maestro User Interface for the sender definition step of creating a job.

It is possible to have more than one address and password configured within the LISTSERV instance and granted the DISTRIBUTE right. Follow the steps above to add additional addresses if desired. Multiple addresses with the DISTRIBUTE right will allow for individual accounts within a group to have individual sender addresses or for distinct groups to use the same LISTSERV instance without sharing other privileges (for example, sharing the rights to send to the same LISTSERV lists).



Tip: Each account group should have a separate distribute account defined in order to prevent inadvertent sharing of LISTSERV List data. For even stricter separation, each group can connect to a different LISTSERV server.

4.1.2 Preparing LISTSERV to Allow Maestro Hosted LISTSERV Lists

On some platforms, it is possible for LISTSERV Maestro to create traditional LISTSERV lists and store the subscription data for these lists such that they are accessible from the LISTSERV Maestro subscription pages. These types of lists are referred to as Hosted LISTSERV Lists or HLLs because their data is "hosted" within LISTSERV Maestro's system database.

First determine whether HLLs are supported on your LISTSERV platform. As far as LISTSERV is concerned, HLLs are simply DBMS-backed lists. DBMS-backed lists are not supported on every LISTSERV platform using every DBMS package.

LISTSERV running on Windows supports any ODBC-compliant database so HLLs are always possible if LISTSERV is running on Windows, even if LISTSERV Maestro is using the internal system database.

If LISTSERV is running on Linux or Solaris, HLLs are possible only if the LISTSERV Maestro system database is defined as an external database using Oracle or DB2. If LISTSERV is running on Mac OS X, HLLs are possible only if the LISTSERV Maestro system database is defined as an external system database using Oracle. If LISTSERV is running on a different operating system (remember that LISTSERV may run on a different platform than the other LISTSERV Maestro components), support for HLLs depends on LISTSERV support for DBMS-backed lists on that platform. The table below summarizes the restrictions for HLL support.

Table 4-1 Restrictions for Hosted LISTSERV List Support

If LISTSERV is running on this Operating System:	LISTSERV Maestro HLL support requirements are:
Windows	Any supported internal or external system database (MySQL, Oracle, DB2, or MS SQL Server).

If LISTSERV is running on this Operating System:	LISTSERV Maestro HLL support requirements are:
Linux Solaris HP-UX AIX	The external system database must be defined using Oracle or DB2.
Mac OS X Tru64 Unix OpenVMS	The external system database must be defined using Oracle.
FreeBSD Linux-S390	No support for Hosted LISTSERV Lists.
VM	LISTSERV Maestro is not supported for use with LISTSERV running on VM.

Assuming the platform supports Hosted LISTSERV Lists, the following steps are necessary to prepare LISTSERV to support Hosted LISTSERV Lists in LISTSERV Maestro.

1. Add a special Postmaster address for use by Maestro



Note: On Windows, if the **Express Setup** option was used to install the full LISTSERV Maestro Suite (including LISTSERV), a postmaster address and password have already been defined for Maestro, and you can skip directly to step 2.

In order to allow LISTSERV Maestro to create Hosted LISTSERV Lists, the LISTSERV Site Administrator needs to register an email address for LISTSERV Maestro under the `POSTMASTER` setting of the LISTSERV Site Configuration file. This should be a separate address from those registered as `DIST_ALLOWED_USERS`, as the `POSTMASTER` address will additionally have the rights to create new lists on the LISTSERV server and to access any LISTSERV list on that server. It is recommended that the `QUIET` and `HIDE` parameters are used for the address, so that it does not receive routine administrative messages from the LISTSERV server. An example from a Windows `SITE.CFG` file might look as follows (while the line may be wrapped in this document, it should not be in the actual configuration file):

```
POSTMASTER=POSTMASTER@EXAMPLE.ORG QUIET:
HIDE:MAESTRODIST@EXAMPLE.ORG
```

2. Prepare the LISTSERV List Archive Folder

If Hosted LISTSERV Lists will be enabled to use LISTSERV's List Archive features in order to keep an archive of postings to LISTSERV lists, then the LISTSERV Site Administrator needs to create a folder on the LISTSERV server in which to store the archive files. This is usually a sub-directory under the main LISTSERV folder (something like `C:\LISTSERV\LISTS` or `/home/listserv/lists.`) This folder location will be later entered into the Administration Hub; see Section 5 [Settings for the Maestro User Interface](#).

3. Prepare LISTSERV for database access

Hosted LISTSERV lists store their subscriber information in a DBMS-back-end, rather than “traditional” LISTSERV list files. In order to use DBMS-driven Hosted LISTSERV Lists, the LISTSERV Site Administrator must configure LISTSERV to access the LISTSERV Maestro system database. First, a compatible database client must be installed and configured on the LISTSERV host. Next, LISTSERV must be enabled to use that client. Finally, LISTSERV must be configured to access the LISTSERV Maestro system database. See the next section for detailed instructions.

4.1.3 Preparing LISTSERV for Database Access

There are two situations in which LISTSERV needs to be prepared for database access:

- If Hosted LISTSERV Lists will be used (see Section 4.1.2 [Preparing LISTSERV to Allow Maestro Hosted LISTSERV Lists](#)), and
- If LISTSERV will be used to retrieve recipients directly from an external user database instead of LISTSERV Maestro retrieving the recipients and passing them on to LISTSERV. In the LISTSERV Maestro User Interface (LUI), this corresponds to the **Let LISTSERV Select Recipients from a Database** recipient type.

Installing and Configuring a Database Client – The System Administrator of the LISTSERV host will need to install and configure a database client on the LISTSERV system. For LISTSERV on Windows, ODBC is the only supported database client (as of the time this document was written). For LISTSERV on Solaris, Linux, AIX, Oracle’s OCI client, and IBM’s DB2 CLI client are the supported clients. For LISTSERV on Mac OS X, Tru64 Unix, and OpenVMS, only Oracle’s OCI client is supported (please consult the LISTSERV documentation for current specific version support). Please consult the vendor documentation for details on installing the appropriate ODBC, OCI or CLI client files for the LISTSERV platform.

Once the database client is installed, a client connection to the LISTSERV Maestro system database needs to be configured. This usually means supplying the client with a host name and port number for your database server, a database name, and a unique service name.

For sites using an external database as the LISTSERV Maestro system database, configure the database client on the LISTSERV server with connection settings for that particular environment. Usually this requires at least a host name for the database server, port number, database name, and unique service name. Consult the vendor documentation for information on configuring these settings.

Details for LISTSERV on Windows Connecting to the Internal System Database for HLL Support – For LISTSERV Maestro installations on Windows using the default internal system database and running LISTSERV on the same host as Maestro, an ODBC driver for MySQL will need to be installed on that host. (At the time this document was written, MySQL provides such a driver on their web site at <http://dev.mysql.com/downloads/connector/odbc/>). Open the ODBC Data Sources control panel, and create a new system DSN using the MySQL driver. The host name configured in the DSN should be `localhost`, the port should be `3306`, and the database name should be `LUI`. Give the DSN a unique name; this name will be assigned to the `ODBC_DSN` parameter in the LISTSERV site configuration file (see below).

For sites using the default internal system database, but with LISTSERV for Windows installed on a different host than LISTSERV Maestro, an ODBC driver for MySQL will again need to be installed on the LISTSERV host. While configuring the ODBC DSN, substitute the actual hostname of the LISTSERV Maestro server for `localhost` (the port number and database name will be the same as above). On the LISTSERV Maestro server, give LISTSERV permission to access the internal system database. Start the MySQL client program `\Program Files\L-Soft\Application Server\lui\database\bin\mysql.exe`, login as root, and then enter the following command:

```
grant all on LUI.* to USERNAME@LISTSERVHOST identified by 'PASSWORD';
```

Replace `USERNAME` with the `ODBC_UID` that you will be assigning to LISTSERV (see below), `LISTSERVHOST` with the name of the host on which LISTSERV is running, and `PASSWORD` with the `ODBC_AUTH` that you will be assigning to LISTSERV (see below).

Enabling LISTSERV To Use a Database Client Driver – On Windows, LISTSERV comes already linked to the Microsoft ODBC system. No additional steps are needed in order to enable DBMS support in LISTSERV (although it is still necessary to configure the ODBC connection and enter specific database information in LISTSERV's Site Configuration file, see below).

On other platforms, if LISTSERV is not already compiled with database support (which is the case by default), then you'll have to re-link the LISTSERV executable to include database support for Oracle or DB2. Make sure to use the current LISTSERV installation kit from <ftp://ftp.lsoft.com/LISTSERV/UNIX>, and follow the instructions in the LISTSERV installation memo for adding DBMS support to the LISTSERV executable.

Configuring LISTSERV to Access the LISTSERV Maestro System Database – For Hosted LISTSERV lists, LISTSERV stores its subscriber data in tables in the LISTSERV Maestro system database. Once there is a database-enabled LISTSERV installation, the LISTSERV Site Administrator next needs to configure LISTSERV to access the LISTSERV Maestro system database. This is done by adding configuration information to LISTSERV's Site Configuration file (`SITE.CFG` or `go.user`). See the LISTSERV Site Manager's Manual at <http://www.lsoft.com/resources/manuals.asp> for detailed information on editing LISTSERV's Site Configuration File). Typically, three parameters are needed: a database identifier, a user name, and a password.

Before adding any of the following parameters to the LISTSERV Site Configuration file, take note of the database server name of the LISTSERV Maestro system database. LISTSERV Maestro automatically generates this server name upon installation. To view the database server name through the Maestro Administration Hub click on the **Global Settings** menu, select **Maestro User Interface**, and then select the **Default LISTSERV Connection** option. Under the **Hosted LISTSERV Lists Settings** section, look for the **Database Server Name** field. The name in this field should begin with `MAESTRODB_`, followed by a string of characters.

Figure 4-1 The Database Server Name

Hosted LISERSV Lists Settings

The settings made here are only required if the group or account using the LISERSV connection configured above is allowed to use Hosted LISERSV Lists.

Usage of List-Archives for Hosted LISERSV Lists:

Archives for Hosted LISERSV Lists allowed, archive location as specified below: ▼

Base folder for archive location (Must already exist at LISERSV host):

D:\Development\MAESTR~1\LISTSERV\LISTS

Database Server Name: **MAESTRODB_elex3jha**
(Must be configured as a database server name at the LISERSV host)

In the example in Figure 4-1, the Database Server Name is MAESTRODB_elex3jha. This server name uniquely identifies the Maestro system database to LISERSV. Enter that database server name into the LISERSV Site Configuration file (SITE.CFG or go.user) exactly as it appears in the Administration Hub – including upper and lower case letters. Utilize the following parameters to configure LISERSV for access to the Maestro system database (depending on the DBMS product in which the Maestro system database is housed). Replace SERVERNAME with the Database Server Name displayed in the Maestro Administration Hub:

- Microsoft ODBC connections (Windows):
 - ODBC_DSN_SERVERNAME – The ODBC_DSN parameter should be the Database Service Name (DSN) that was already configured in the Windows ODBC control panel (see above for DSN configuration information).
 - ODBC_UID_SERVERNAME – The ODBC_UID parameter should contain the user name under which LISERSV should connect to the DSN configured above. This user name (and corresponding password) should already be registered in the DBMS. If the internal system database installed with LISERSV Maestro is being used, and LISERSV is running on the same host, the user name for the LUI database is luiuser.
 - ODBC_AUTH_SERVERNAME – The ODBC_AUTH parameter should contain the password that corresponds to the ODBC_UID configured above. If the internal system database installed with LISERSV Maestro is being used and LISERSV is running on the same host, the password for the luiuser user is lui.
 - Example for Windows/ODBC:


```
ODBC_DSN_MAESTRODB_dtn7nzn3=LUI
ODBC_UID_MAESTRODB_dtn7nzn3=luiuser
ODBC_AUTH_MAESTRODB_dtn7nzn3=lui
```
- Oracle OCI connections (UNIX/Linux):
 - OCI_CONNECT_SERVERNAME – The OCI_CONNECT parameter should contain an OCI connect string (typically configured in a TNSNAMES.ORA file. Consult the Oracle documentation for connect string configuration.)

- `OCI_UID_SERVERNAME` – The `OCI_UID` parameter specifies the user name under which LISTSERV will connect to the database specified by the connect string. This user name (and corresponding password) should already be registered in the DBMS. (If a user name and password have already been specified as part of the `OCI_CONNECT` parameter, then the `OCI_UID` and `OCI_PWD` parameters may not be necessary.)
- `OCI_PWD_SERVERNAME` – The `OCI_PWD` parameter specifies the password to be used for the `OCI_UID` configured above.
- Example for UNIX/OCI (with UNIX, each parameter must be exported):

```
OCI_CONNECT_MAESTRODB_dtn7nzn3="MYDATA"
OCI_UID_MAESTRODB_dtn7nzn3="MYUSER"
OCI_PWD_MAESTRODB_dtn7nzn3="MYPASS"
```

- IBM DB2 connections (UNIX/Linux):
 - `CLI_DSN_SERVERNAME` – As with the `ODBC_DSN`, the `CLI_DSN` specifies some Database Service Name (DSN) already configured in the CLI configuration.
 - `CLI_UID_SERVERNAME` – The `CLI_UID` parameter specifies the user name under which LISTSERV should connect to the `CLI_DSN` configured above. This user name (and corresponding password) should already be registered in the DBMS.
 - `CLI_AUTH_SERVERNAME` – The `CLI_AUTH` parameter supplies the password for the `CLI_UID` configured above.
 - Example for UNIX/CLI (with UNIX, each parameter must be exported):

```
CLI_DSN_MAESTRODB_dtn7nzn3="MYDATA"
CLI_UID_MAESTRODB_dtn7nzn3="MYUSER"
CLI_AUTH_MAESTRODB_dtn7nzn3="MYPASS"
```

In addition to the ODBC, OCI, and CLI parameters listed above, LISTSERV must also be supplied with a default database connection (without the additional `SERVERNAME` parameter). This default connection may contain invalid “dummy” settings, but it must be present in order for LISTSERV to enable its database functions. LISTSERV looks in its site configuration file for a parameter called `ODBC_DSN`, `OCI_CONNECT` or `CLI_DSN` (without the additional `SERVERNAME`) in order to determine whether the database extensions should be initialized. Enter “dummy” information in addition to the actual Maestro system database information, as in the example below. The following example is for OCI; substitute ODBC or CLI parameter names as appropriate:

```
OCI_CONNECT="DUMMY"
OCI_UID="DUMMY_USER"
OCI_PWD="DUMMY_PASS"
OCI_CONNECT_MAESTRODB_dtn7nzn3="MAESTRODATA"
OCI_UID_MAESTRODB_dtn7nzn3="SCOTT"
OCI_PWD_MAESTRODB_dtn7nzn3="TIGER"
```

In the example above, the `OCI_CONNECT`, `OCI_UID` and `OCI_PWD` parameters define the “dummy” database connection, and the `OCI_CONNECT_MAESTRODB_dtn7nzn3`, `OCI_UID_MAESTRODB_dtn7nzn3` and `OCI_PWD_MAESTRODB_dtn7nzn3` parameters define the actual connection settings for the Maestro system database.



Important: If there is not a default DBMS connection defined, (`ODBC_DSN`, `OCI_CONNECT`, or `CLI_DSN`) then LISTSERV’s DBMS features will be disabled and Hosted LISTSERV lists will not work properly, even if there is an `ODBC_DSN_SERVERNAME`, `OCI_CONNECT_SERVERNAME`, or `CLI_DSN_SERVERNAME` parameter defined. It is permissible to use a “dummy” setting (one which contains invalid “placeholder” settings) for the default DBMS connection, but a default connection must be configured in order for LISTSERV’s DBMS support to be enabled.

4.2 Preparing LISTSERV Maestro to Send DISTRIBUTE Jobs to LISTSERV

After LISTSERV has been configured for use with LISTSERV Maestro, the LISTSERV configuration information must be entered into LISTSERV Maestro's Administration Hub. The LISTSERV connection can be configured at several levels. The widest level of setting is the Global Component Setting, which defines the global default LISTSERV connection on the application level (see Section 5 [Settings for the Maestro User Interface](#)). These settings will be used as defaults for all accounts that do not have individual settings for group or user level.

The next level is the default setting for a group. The group default LISTSERV connection may connect to a different instance of LISTSERV, or may use different settings than the global default. To set the LISTSERV settings for an entire group in the user list, click on the group name (appears only if the account belongs to a group). If defined on the group level, the group settings will override the global default settings for that group. The settings will apply to all accounts in the group, except for the **Client Address for Mail Delivery** settings (**Address** and **Password**). These two settings may be configured individually for each account in the group, provided LISTSERV is configured with additional addresses with the DISTRIBUTE right. Accounts in the group for which they are not configured will use the group's settings as the default.

Individual accounts are the last level of settings for the LISTSERV Connection. Individual accounts can belong to a group (group user) or not belong to a group (single user). Settings for individual accounts will override default group and global settings. To configure the LISTSERV connection for an individual account, click on the user name from the user list. The screen that opens is different depending on whether the account is a group user or a single user.

Single user accounts can use a separate LISTSERV instance and would therefore need to have all LISTSERV connection settings defined. Or, single user accounts may use a different client address and password than the global default. Group user accounts can only define the **Client Address for Mail Delivery** settings. These individual group settings result in a different LISTSERV email address for each user so that jobs can be identified by owner in the LISTSERV logs.

Properly specifying the LISTSERV host name settings, found on the LISTSERV Connection screens, is another important aspect to preparing LISTSERV Maestro to process DISTRIBUTE jobs from LISTSERV. This is important for three reasons:

- **LISTSERV Maestro to LISTSERV system communication** – LISTSERV Maestro communicates with LISTSERV using the TCPGUI port. For this reason, LISTSERV Maestro needs to know the name of the server where LISTSERV is running. Enter this server name in the **LISTSERV Host** field.
- **Host Name For Return Path** – When email is sent out over the Internet, the return path, which allows undeliverable messages to be returned to the sender (also known as the RFC 821 address), must include a known external name, otherwise the bounced mail cannot be returned and LISTSERV Maestro cannot automatically process and report on bounces. Use the **External Host Name** drop-down menu to help identify this external name.
- **Host Name For List Addresses** – When doing mailings that are based on a normal LISTSERV list, the address of the list in use must include a known external name of the server hosting the LISTSERV instance for list communication to function correctly (something like `listname@hostname`). Use the **External Host Name** drop-down menu to help identify this external name.

Usually, a server is only given a single host name by which it can be reached from other computers, including both internal intranet computers and external computers on the Internet. In this case, enter that name into the **LISTSERV Host** field of the LISTSERV Connection screen. Click the **External Host Name** drop-down menu and select **Use LISTSERV host name as given above**.

When the LISTSERV server is given several different names (or appears to have several names) a different setup is required. These situations often stem from optimizing a high performance installation of LISTSERV Maestro. In this case, follow the instructions in Section 4.2.1 [Specifying the LISTSERV Host with Different Internal and External Names](#).

For very high volumes, it may be desirable to have a separate LISTSERV installation solely for the purpose of processing bounces. In that case, follow the instructions in Section 4.2.2 [Specifying a Separate LISTSERV Instance for Processing Bounces](#).

4.2.1 Specifying the LISTSERV Host with Different Internal and External Names

A common optimizing set up is to have LISTSERV on one server inside a firewall with only an internal name, and the SMTP service on another server outside the firewall with an external name. With this set up, LISTSERV, installed on the server with the internally known name, is visible by this name to inside users. For all outside purposes, such as the return path and list email addresses, LISTSERV “appears” to actually be running on the SMTP server with the external name. This is because the SMTP service is connected to the actual LISTSERV instance on the internal server. When viewed remotely, LISTSERV appears to have two host names: one internally known and one externally known.

To enter a separate external host name, select **Use special external LISTSERV host name specified below** from the drop-down menu. An edit box will appear. Enter the host name in the edit box.

Figure 4-2 Using A Different External Host Name For LISTSERV

LISTSERV Connection

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

LISTSERV Host: <default undefined>

External Host Name: <default undefined>

LISTSERV TCPGUI Port: Default: 2306

Client Address for Mail Delivery and Bounce Reporting (Must have DISTRIBUTE Right):

Address: <default undefined>

Password: <default undefined>

Client Address for LISTSERV Server Administration (Must be POSTMASTER):

Address: <default undefined>

Password: <default undefined>

Use the following rules for entering information on the LISTSERV Connection screen:

- For the internal communication, (the **LISTSERV Host** field) always specify a host name that points to the server where LISTSERV is actually running, not to the instance where LISTSERV only “appears” to be running, when in reality it is SMTP running on that server. Also, the LISTSERV Maestro server must be able to resolve that host name to the actual IP address of that server.
- For the **External Host Name** field, always specify the host name of LISTSERV as seen by outside clients (the Internet), even if that name is actually only an alias for the host or if it points to a server where only the SMTP instance is running.

4.2.2 Specifying a Separate LISTSERV Instance for Processing Bounces

For very large volume installations of LISTSERV Maestro, a separate instance of LISTSERV can be used just to process bounces. From the LISTSERV Connection screen, select the **Use dedicated server** option if you want more settings for this LISTSERV instance to appear.

If the dedicated bounce processing host has only a single name, enter that name into the **LISTSERV Host** field of the LISTSERV Connection screen. Select **Use LISTSERV host name as given above** from the **External Host Name** drop-down menu. Fill in the other appropriate information, the TCPGUI port, the client address and password following the same rules outlined in Section 4.2 [Preparing LISTSERV Maestro to Send DISTRIBUTE Jobs to LISTSERV](#).

If the dedicated bounce-processing host has more than one name (or appears to), then select **Use special external LISTSERV host name specified below** from the drop-down menu. An edit box will appear. Enter the host name in the edit box.

4.3 Using Existing Lists with LISTSERV Maestro

LISTSERV Maestro can provide access to existing LISTSERV lists, presenting a drop-down menu of available lists on the Source page of the Recipient Definition wizard when **Send to an Existing LISTSERV List** is selected on the Options page. In order to do this, follow these instructions:

For each existing LISTSERV list to be added to the drop-down menu, insert a line in the list header containing the keyword "Owner=". Add the email address that was configured with the DISTRIBUTE right (LISTSERV client address) to the right side of the "=" sign. This can be accomplished using email for all LISTSERV servers. If LISTSERV's Web Interface is installed and configured, it can be used to edit the list header. See the LISTSERV List Owner's Manual for detailed information on editing LISTSERV list headers.

For example, in the sample list header below, the lines highlighted in gray have been added. The address to the right of the "=" sign is the address that has the DISTRIBUTE right in the LISTSERV instance where this list is located. The line before that sets this owner to "quiet" meaning that no mail will ever be sent to that address. This is useful in case that address does not resolve into a real mailbox, but only exists to allow DISTRIBUTE jobs through LISTSERV Maestro. The example below shows how the list header might look when edited through the LISTSERV web interface; if edited through email, each header line would be preceded by an asterisk (*):

Women's Club

```
Notify= Yes
Editor= user@example.org
Owner=someone@example.org
Owner=quiet:
Owner= maestro@companyserver.example.org
Moderator= All
Sizelim= 1M
Subscription= By Owner
Subscription= Confirm
Ack= Yes
Confidential= Yes
Validate= No
Reply-to= Sender,Respect
Send= Private
Errors-To= Owner
Notebook= Yes,E:\LISTS\WOMENS_CLUB,Weekly,Private
```

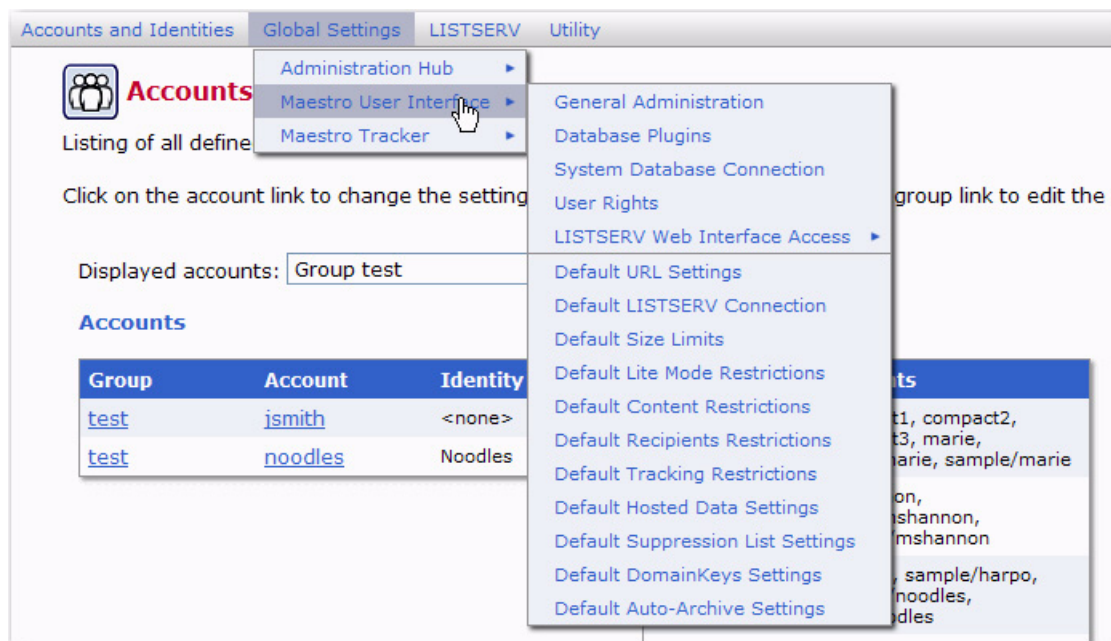
For institutions that have many lists, it is likely that different people will need to have access to different lists. If this is the case, people that work on the same list or set of lists will have to be placed in the same group. The group can have its own default LISTSERV connection and/or LISTSERV client address and password. If different addresses are assigned to individuals in a group (multiple LISTSERV client addresses), then all those addresses must appear in the list headers as `Owner=`.

Section 5 Settings for the Maestro User Interface



To select settings for the Maestro User Interface, click on the **Global Settings** menu, and select the **Maestro User Interface** option. The top selections define application settings, and the bottom selections define application defaults. Click on an option to define or edit the settings.

Figure 5-1 Global Settings for the Maestro User Interface



5.1 Application Settings

Application settings are general settings that affect LISTSERV Maestro globally. The Maestro Administrator may change them.

5.1.1 General Administration

The **General Administration** setting defines general Maestro User Interface settings. These settings include:

- **Backup folder** – Defines the folder where the daily backup of the Maestro User Interface is written. If left empty, the default backup folder will be used. Use a relative or absolute path. Relative paths are relative to the Maestro User Interface’s home folder.
- **Event transfer interval** – Tracking events are initially collected in Maestro Tracker, but before they become available for reports, they need to be transferred to the Maestro User Interface. To prevent the components from being overburdened, these transfers happen in “bursts,” and this parameter defines the time interval between bursts. As a result, there will not be any apparent changes in reports until

the next interval has passed, transferring more tracking data from Maestro Tracker to the Maestro User Interface.



- **Job archive folder** – Defines the folder where archived jobs are saved. Archived jobs are special ZIP archive files that are removed by the administrator from the Delivered Jobs listing in the Maestro User Interface. If left empty, the folder named “archive” inside the Maestro User Interface application home folder will be used.
- **Time for daily system metrics update** – Defines the time for your daily system metrics to update. The default is 12:00.
- **Send a mail with the system metrics after each update** – Check if you’d like to have an email message sent to show you how much disk space is being used.

Figure 5-2 General Administration of the Maestro User Interface - General Settings

General Settings

Backup folder:
(Leave empty for default: "backup".)

Event transfer interval (from Maestro Tracker): min.

Job archive folder:
(Leave empty for default: "archive".)

Time for daily system metrics update: (hh:mm)
(This time is relative to the time zone of the server)

Send a mail with the system metrics after each update.
(Administrative E-mail Notifications must be enabled.)

- **Advanced Security Options** – These settings define the number of allowed unsuccessful login attempts a user can have before being locked out of the LISTSERV Maestro User Interface. Using the [**Unlock all currently locked accounts**], the administrator can unlock all accounts if a user needs to access the User Interface before the “lock login time” has expired.



Note: This can also be enabled for the HUB. For more information, see Section 17.3.1 [Securing the Administration Hub](#).

Figure 5-3 General Administration of the Maestro User Interface - Advanced Security Options

Advanced Security Options

Limiting the number of allowed unsuccessful login attempts is a measure to avoid external automated dictionary attacks. Supply the maximum number of allowed unsuccessful login attempts below to activate this advanced protection.

Max No. of Unsuccessfull Logins: (leave empty for unlimited number)

Lock Login Time in Minutes:

- **Runtime Administration** – These settings allow the administrator to influence the availability of the Maestro User Interface – for example, in the event of a system shutdown. The administrator can disable the Outbox, lock login access, present a message at the top of each screen to logged in users while the login is locked, and create a message that appears to any user trying to login while the login is locked.



Tip: Decide on a time slot each week to perform non-emergency maintenance of the server (for example, software upgrades). Let users know about this in advance so they know to avoid sending jobs right before that time slot. When taking the system down for non-emergency maintenance, disable the Outbox and login access an hour ahead (or whatever time seems appropriate) to give the users time to finish their current activities and log out in time.

Figure 5-4 General Administration of the Maestro User Interface - Runtime Administration

Runtime Administration

Multiple Logins: Disallow multiple logins with the same user account.

Outbox: Sending is disabled.

Login Access: LISTSERV Maestro User Interface is locked.

Message that will be shown instead of login page while login is locked:

Message that will be shown at top of each page while login is locked:

- **Subscriber Access Page Translations** – Translations for the Subscriber Pages are maintained externally within the L-Soft Resource Translation Tool and are made available to LISTSERV Maestro as JAR files. See Section 6.5 [Refreshing the Subscriber Page Translations](#) for more information.

5.1.2 Database Plugins

The **Database Plugins** setting registers and unregisters database plugins. Database plugins allow LISTSERV Maestro to communicate with databases. For more information, see Section 9.2 [Registering a Database Plugin](#).

Figure 5-5 Database Plugins

Database Plugins

The list below shows all registered database plugins.

Click on the "unregister" link to unregister it.
Use the "Register New Database Plugin" button to register a new plugin.

Plugin Name	Full Class Name
SQL Server jTDS Driver Database Plugin	com.lsoft.lui.db.sqlserver.JTDSDriverPlugin Unregister

5.1.3 System Database Connection

The **System Database Connection** setting defines the settings for the Maestro System Database. (For more information on the system database, see Section 10 [The System Database](#).) These settings include:

- **Maximum number of buffered connections** – Defines the maximum number of “open” database connections the Maestro User Interface will keep open at any time. After the Maestro User Interface has finished using a connection, it will not close the connection, but keep it open as a buffered open connection.
- **System Database connection choice** – Select the **Use the internal database as the System Database** option to use the internal database (based on MySQL) as the system database. Select the **The following external database is used as the System Database** option to use an external database as the system database. In this case, select the corresponding database plugin from the drop-down menu. Once a plugin has been selected, a set of input fields will appear. Enter details for the database connection such as server name, database port, database name, user name and password. The exact details depend on the plugin selected.

Figure 5-6 System Database Connection

System Database Connection

Maestro System Database Connection

On this page you define the settings of the System Database that the Maestro User Interface uses to store its internal system data in.

Maximum number of buffered connections:

Use the internal database as the System Database

The following external database is used as the System Database:

Database Plugin:

Connection Details

Database Name:

SQL Server User Name:

Password:

Database Host Name:

TCP/IP Port or Instance Name:

(Note: Changes in this category require a restart of the Maestro User Interface to take effect.)



5.1.4 User Rights

The **User Rights** setting defines the user rights for corresponding accounts. For more information, see Section 7.3 [Managing User Rights](#).

Figure 5-7 User Rights

User Rights Management for Maestro User Interface

Use the checkboxes to define the user rights of the corresponding accounts.

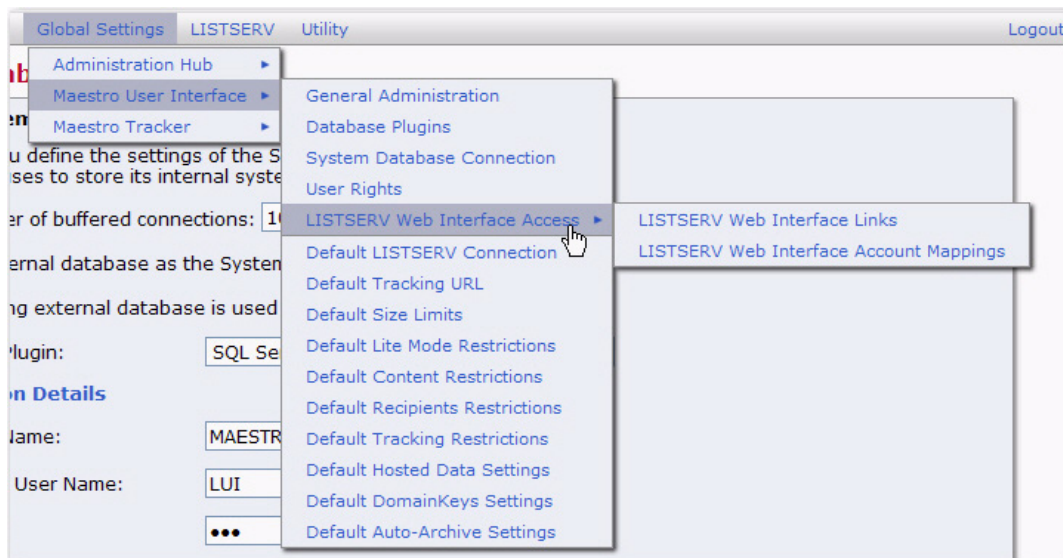
Displayed accounts:

User	Create Jobs	Create Reports	Admin Sender Profiles	Admin Drop-Ins	Admin Content Templates	Admin Target Groups	Admin Datasets	Link Datasets & WA	Designated Job Owner
bparker	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<This Account>
francoise	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	jharlan
jharlan	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<This Account>
liam	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<This Account>
mshannon	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	jharlan

5.1.5 LISTSERV Web Interface Access

The **LISTSERV Web Interface Access** setting configures if and how direct access between LISTSERV Maestro and the LISTSERV Web Interface of the connected LISTSERV instance will be possible. For more information, see Section 22 [LISTSERV & LISTSERV Maestro Integration](#).

Figure 5-8 LISTSERV Web Interface Access



5.2 Application Default Settings

Application default settings are used to set system-wide defaults. LISTSERV Maestro will use default settings if no other settings have been entered at the group or user level. To use default settings, leave all other settings at the group and user level blank. If different settings are entered at the group or user level, they will override the default settings.



5.2.1 Default URL Settings

The **Default URL Settings** define the settings for the various URLs used by the User Interface, and these settings will be used as defaults for all accounts that do not have individual settings on the group or user level.

Figure 5-9 Default URL Settings

URL Settings

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

Default URLs for Users

Default LUI Access URL for Users
This LUI access URL is used by normal LUI users to access the LUI user interface.

Default: http://<Host name of server running the LUI component>

Default LUI Access URL Aliases for Users
These URLs are aliases for the LUI access URL.
If the user accesses any of the alias urls, he is redirected to the official URL defined above.

Default: <none>

Default HUB Access URL for Users
The HUB access URL is needed by the system whenever a normal user changes his password.
This URL is not used by the users directly.

Default: http://<Host name of server running the HUB component>

Default Tracking URL
The tracking URL is used to generate the tracked links in an e-mail.
Tracking is disabled if no tracking URL is specified, either directly or by default!

Default: <none>

URLs for Admin
By default, the admin also uses the URLs specified above, unless they are overridden here:

LUI Access URL for Admin
This LUI access URL is used by the admin to access the LUI user interface.
Important: The admin must specify exactly the protocol, hostname and port defined here to be allowed to login to LUI.

Default: http://<Host name of server running the LUI component>

HUB Access URL for Admin
The HUB access URL is used by the admin to switch from the LUI to the HUB user interface or to login to the HUB directly (if necessary).

Default: http://<Host name of server running the HUB component>

OK Cancel

If defined on the application level, then these settings will be used as defaults for all accounts that do not have individual settings on the group or user level.

- **Default LUI Access URL for Users**

This LUI access URL is used by the LUI users, to access the LUI user interface. For some LUI interface pages, the system needs to include this access URL on the page itself (for example to create a link). The URL that the system uses for this is generated using the setting supplied here. If not supplied, the displayed default URL will be used.

Specify the access URL with the full protocol and hostname, optionally followed by the port. For example,

```
http://luihost.domain.com or https://secureluihost.domain.com:1443
```



Note: A user will only be allowed to login to LUI if in the URL he uses to access LUI, he uses the exact same protocol and hostname (and optionally port) as defined here. If the LUI server has several host names assigned, then using a different host name may successfully display the initial login page, but the login of the user will not be accepted (the system will behave as if the user account does not exist or the wrong password was specified).

- **Default LUI Access URL Aliases for Users**

The LUI Access URL Aliases will be the aliases for the LUI Access URL. If a user accesses LUI with any of the alias urls, then their login is accepted, and after login, they will be automatically redirected to the official LUI Access URL.

Specify a list of alias URLs (one per row), each with the full protocol and hostname, optionally followed by the port. For example,

```
http://luihostalias.domain.com
```

or

```
https://secureluihostalias.domain.com:1443
```

- **Default HUB Access URL for Users**

This HUB access URL is used behind the scenes to allow the system to display the password change page for users. This URL is not used by users directly. To display the password change page, the system needs to be able to generate a URL that correctly points to the HUB user interface. The URL that the system uses for this is generated using the setting supplied here. If not supplied, the displayed default URL will be used.

Specify the access URL with the full protocol and hostname, optionally followed by the port. For example,

```
http://hubhost.domain.com or https://securehubhost.domain.com:1443
```

- **Default Tracking URL**

The tracking URL is used to generate the tracked links in the email messages that have tracking enabled. The URL that the system uses for this is generated using the setting supplied here. If not supplied, the displayed default URL will be used.

Specify the tracking URL with the full protocol and hostname, optionally followed by the port. For example,

`http://trkhost.domain.com` or `https://securetrkhost.domain.com:1443`



Important: There is no default for the tracking URL. Therefore, as long as no tracking URL is defined for a given account or group, then tracking is disabled for this account/group.



Tip: For more details on tracking URLs, see Section 5.3 [Setting the Default Tracking URL](#).

- **Default LUI Access URL for Admin**

This LUI access URL is used by the admin to access the LUI user interface (and through it, also the HUB). For some LUI interface pages, the system needs to include this access URL on the page itself (for example, to create a link). The URL that the system uses for this is generated using the setting supplied here. If not supplied, the displayed default URL will be used.

Specify the access URL with the full protocol and hostname, optionally followed by the port. For example,

`http://luihost.domain.com` or `https://secureluihost.domain.com:1443`



Note: The admin will only be allowed to login to LUI if, in the URL he uses to access LUI, he uses the exact same protocol and hostname (and optionally port) as defined here. If the LUI server has several host names assigned, then using a different host name may successfully display the initial login page, but login will not be accepted (the system will behave as if the wrong password was supplied).

- **HUB Access URL for Admin**

This HUB access URL is used behind the scenes to allow the admin to switch from the LUI user interface to the HUB user interface. In certain cases, it may also be used to access the HUB user interface directly. The URL that the system uses for this is generated using the setting supplied here. If not supplied, the displayed default URL will be used.

Specify the access URL with the full protocol and hostname, optionally followed by the port. For example,

`http://hubhost.domain.com` or `https://securehubhost.domain.com:1443`

Click **[OK]** to submit and save any changes, or **[Cancel]** to exit the page without saving your changes.

5.2.2 Default LISTSERV Connection

The **Default LISTSERV Connection** setting defines the default LISTSERV connection, the default settings for LISTSERV Hosted Lists, and the LISTSERV Connection for automatic bounce handling. The default setting is used for all accounts that do not have single user or group LISTSERV connections defined. LISTSERV settings defined at the user or group level will override the default settings for only those users or groups. As a result, it is possible to have some users and groups using the default LISTSERV settings and other users and groups using settings defined at the user or group level.

Figure 5-10 Default LISTSERV Connection



Note: If there will be several unrelated groups, it is recommended not to define default settings for client addresses and passwords, as each group should use a separate address, defined in the group's LISTSERV Connection settings. If each group will access a separate LISTSERV instance, do not specify any default LISTSERV connection information.

- **LISTSERV Host** – Enter the host name of the server that is actually running LISTSERV. LISTSERV Maestro will use this host name to look up the server running LISTSERV and connect to it using the TCPGUI port. Do not use a server name or alias that only appears to the outside clients to be running LISTSERV, such as the SMTP server name. For more information, see Section 4.2.1 [Specifying the LISTSERV Host with Different Internal and External Names](#).
- **External Host Name** – If different from the LISTSERV Host, enter the host name of the server running LISTSERV as seen by outside clients such as Internet. This host name can be an alias or point to the SMTP server. For more information, see Section 4.2.1 [Specifying the LISTSERV Host with Different Internal and External Names](#).
- **LISTSERV TCPGUI Port** – Enter the port number on the LISTSERV host where LISTSERV listens for TCPGUI connections. The default is 2306.
- **SMTP Worker Pool** – Lets you specify a LISTSERV worker pool to use for specific delivery situations. You can specify different worker pools for standard deliveries and for test deliveries. Worker pools are an advanced LISTSERV feature, see the LISTSERV documentation for details. Leave the field empty to use the displayed default. The system default is "no worker pool", i.e. normal distribution by LISTSERV. If you do not know what worker pools are or what they are used for, you should always stick to this system default.

- **Client Address for Mail Delivery and Bounce Processing** – If a dedicated bounce server is used, then this entry reads “Client Address for Mail Delivery”.
 - **Address** – Enter the email address that has been configured in LISTSERV to have the right to send DISTRIBUTE jobs. See Section 4.1 [Preparing LISTSERV for LISTSERV Maestro](#).



Notes: This is not the address that will be used as the “From” address of the actual email messages sent to the recipients; that type of address is defined for each individual email job during the creating of the job.

This address must be configured at the LISTSERV instance to have the right to DISTRIBUTE jobs. To do this, the LISTSERV configuration file must be set up to grant the right to send DISTRIBUTE jobs to this address (it must be listed in one of the POSTMASTER or DIST_ALLOWED_USERS configuration parameters). For more information, see the LISTSERV documentation.

- **Password** – Enter the password association in LISTSERV with the DISTRIBUTE address specified above. See Section 4.1 [Preparing LISTSERV for LISTSERV Maestro](#).
- **Client Address for LISTSERV Server Administration**
 - **Address** – This should be the email address that was added as a POSTMASTER in the LISTSERV Site Configuration file.
 - **Password** – Enter the LISTSERV password that was registered for the Client Address entered above.

If Hosted LISTSERV Lists will be used, the settings for those lists need to be configured as well.

- **Usage of List-Archives for Hosted LISTSERV Lists** – If postings sent to the LISTSERV lists will not be archived, then select **No archives allowed for LISTSERV Lists**. If there will be archived list postings, then select the option to allow archives for Hosted LISTSERV lists, and supply the archive folder location that was configured in section 4.1.2 [Preparing LISTSERV to Allow Maestro Hosted LISTSERV Lists](#).

Figure 5-11 Default LISTSERV Connection - Hosted LISTSERV Lists Settings

Hosted LISTSERV Lists Settings

The settings made here are only required if the group or account using the LISTSERV connection configured above is allowed to use Hosted LISTSERV Lists.

Usage of List-Archives for Hosted LISTSERV Lists:

Archives for Hosted LISTSERV Lists allowed, archive location as specified below:

Base folder for archive location (Must already exist at LISTSERV host):

Database Server Name: **MAESTRODB_elex3jha**
(Must be configured as a database server name at the LISTSERV host)

In high volume environments, a special LISTSERV instance that is dedicated to handling bounced mail may be used. If this is the case, select **Use dedicated server** and then

define the settings of this second LISTSERV instance in the lower fields. For more information, see Section 4 [LISTSERV and LISTSERV Maestro](#).

Figure 5-12 Default LISTSERV Connection - Automatic Bounce Handling

LISTSERV Connection for Automatic Bounce Handling

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

No dedicated bounce server
(use standard LISTSERV server, see above)

Use dedicated server
(settings to be supplied when this option is selected)

5.2.3 Default Size Limits

The **Default Size Limits** setting sets a size limit for email messages and any file uploaded to the system. The size limit for an email message applies to the total byte size of the message (after all transfer encoding and MIME multipart wrappers have been applied). If the message exceeds the size limit, the delivery will fail. The size limit for all uploaded files includes recipient lists, attachments, image files, and so on.

Figure 5-13 Default Size Limits

Size Limits

Size Limits

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.
(Values without a suffix are interpreted as "Bytes". To specify a value in "Kilobytes" or "Megabytes", append a "K" or "M", respectively. (Samples: "150K" or "5M").

Maximum Message Size: Default: No Limit.

Maximum File-Upload Size: Default: 50M

5.2.4 Default Content Restrictions

The **Default Content Restrictions** setting defines AOL Rich Text settings for an alternative part of an HTML message. (The AOL Rich Text setting is obsolete and not recommended except in special cases.) Create a set of parameters to set up a list of files or URLs that are available to use as drop-in content elements; see the online help for information.

Figure 5-14 Default Content Restrictions

Content Restrictions

Content Restrictions

Allow usage of AOL format alternative in HTML mail

Do not allow usage of AOL format alternative in HTML mail

Drop-In Content Restrictions

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

Security Issue: Drop-In content elements of type "file" and "URL" will access the files or URLs entered by the user in the context of the **server**. Therefore, in order to protect sensitive or other non-public information, you need to designate specific files, folders, and URLs that users will be able to access. See the help page for more information.

Prefix-Strings of files to which access is allowed:

Default: No file access allowed

Prefix-Strings of URLs to which access is allowed:

Default: No URL access allowed

OK Cancel

5.2.5 Default Recipients Restrictions

The Default Recipients Restrictions screen is split into two sections. The top section, **Recipients Type Restrictions**, sets the type of recipients the user is able to use for an email job, in addition to mail limits. If **Disabled** is selected, then that option will appear grayed out in the Maestro User Interface, and the user will not be able to select it. If **Hidden** is selected, the recipient type will be disabled and will not appear at all in the Maestro User Interface.

To limit the number of email jobs a group or user sends per day and/or month, click the **Enable** option next to any of the permissible recipient types (see list below). Next, in the **Mail Limits** fields, enter the per day and per month limits. This limit is used to prevent excessive limits for those wishing to control the amount of email being sent (e.g. for server limitation reasons or for license limitations). Limiting the number email jobs sent can only be defined for the following recipient types:

- **Upload a Recipients Text File** (Limits can be defined for both Recipients and Target Groups.)
- **Select Recipients from a Database** (Limits can be defined for both Recipients and Target Groups.)
- **Determine Recipients Based on Reaction on Another Job** (Limits can only be defined for Recipients.)

- **Sent to a Hosted Recipient List** (Limits can only be defined for Target Groups; in addition, they are automatically enabled if a HRL already exists.)



Note: Recipient types that are controlled by LISTSERV (**Send to an Existing LISTSERV List** and **Let LISTSERV Select from a Database or LDAP Directory**) cannot have email limitations applied to them.

Figure 5-15 Default Recipients Restrictions - Recipients Type Restrictions

Recipients Restrictions

Recipients Type Restrictions

Define which recipients types are available by default.
 Select "Enabled" to enable a recipients type.
 Select "Disabled" to disable but still display a certain recipients type.
 Select "Hidden" to disable and hide a certain recipients type.

Recipients types managed by LISTSERV Maestro	Enabled	Disabled	Hidden
Upload a Recipients Text File (Recipients and Target Groups)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Select Recipients from a Database (Recipients and Target Groups)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Determine Recipients Based on Reaction on Another Job (Recipients only)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Send to a Hosted Recipient List (Target Groups only)	(automatically enabled if a HRL is present)		

Supply the mail limits for the recipients types above. Leave the fields empty to inherit the default values, supply 0 to use unlimited mail delivery.

Mail Limits: per day (Default: unlimited), per month (Default: unlimited)

Recipients types managed by LISTSERV	Enabled	Disabled	Hidden
Send to an Existing LISTSERV List (Recipients and Target Groups)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Let LISTSERV Select Recipients from a Database or LDAP Directory (Recipients and Target Groups)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

Note: No mail limit is applicable for the recipients types managed by LISTSERV.

Other	Enabled	Disabled	Hidden
Send to a Recipient Target Group (For the recipients types enabled above)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

Note: The mail limit for target groups depends on the underlying recipients type.

The bottom section, **Recipients Upload Restrictions**, contains a text box for the administrator to enter in allowable paths or path prefixes leading to files on a server accessible to the Maestro User Interface. These files are used for uploading "just before sending" CSV files for recipients definitions. If left blank, CSV files used for "just before sending" recipient definitions in the recipients definition wizard will not be allowed.

Figure 5-16 Default Recipients Restrictions - Recipients Upload Restrictions

5.2.6 Default Tracking Restrictions

The **Default Tracking Restrictions** setting defines the type of tracking that will be available, by default. To enable a tracking type, select **Enable**. To disable, but still display the tracking type, select **Disabled**. To disable and hide a tracking type, select **Hidden**.



To enable Action Tracking for the entire User Interface, select **Enable Action Tracking**. To disable, select **Disable Action Tracking**.

Figure 5-17 Default Tracking Restrictions

	Enabled	Disabled	Hidden
Personal Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Anonymous Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Unique Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Blind Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

5.2.7 Default Hosted Data Settings

The **Default Hosted Data Settings** define the settings for hosted recipient data. These settings include:



- **General Hosted Data Settings** – Defines whether or not all datasets are closed for maintenance. This will globally close all datasets for maintenance (in preparation of a server upgrade, etc.). This setting overrides all individual dataset settings. On the dataset overview screen, the text “All datasets closed for maintenance by system administrator” will be displayed (so the data administrator does not wonder why all of the datasets are closed).

Figure 5-18 Default Hosted Data Settings - General Settings

General Hosted Data Settings

Define if **all** datasets shall be marked as "closed for maintenance", which disallows subscriber access.

All datasets are closed for maintenance.

- **Hosted List Settings** – Select whether or not users/groups may create Hosted Recipient Lists or Hosted LISTSERV Lists.

Figure 5-19 Default Hosted Data Settings - Hosted List Settings

Hosted List Settings

Define which types of Hosted Lists may be created by the users/groups by default:

Users/groups may create Hosted Recipient Lists

Users/groups may create Hosted LISTSERV Lists

Note: If Hosted LISTSERV Lists are enabled, then the settings in the sub-section "Connection Settings for Hosted LISTSERV Lists" of the LISTSERV Connection of all groups and accounts to which this applies must also be configured.

If Hosted LISTSERV Lists creation is enabled above or by individual user/group settings, define the name suffix settings for LISTSERV lists.

Do not use name suffixes for LISTSERV lists

Use automatically generated name suffixes for LISTSERV lists

Use individual user/group name suffixes for LISTSERV lists
(Only groups or users with a specified list name suffix will be allowed to created hosted LISTSERV lists.)

If Hosted LISTSERV Lists will be used, optional suffixes for the LISTSERV lists can be used by selecting the option. Since there may not be more than one list with the same name on the same LISTSERV server, employing automatically generated or individual group/name suffixes for LISTSERV list names prevents list creation errors due to duplicate list names. See the online help for more details about name suffixes.



Tip: By installing separate LISTSERV instances, and configuring the LISTSERV connection for each user/group to use a different LISTSERV instance, then the use of suffixes is not necessary to distinguish between lists. Separate instances will also give each group a separate LISTSERV Web Interface for their lists.



- **Hosted Data Import Restrictions** – Define the file import restrictions and prefix strings for the files that allows access. Access to the server files is restricted and only files that start with one of the prefixes defined here are allowed.

Figure 5-20 Default Hosted Data Settings - Hosted Data Import Restrictions

Hosted Data Import Restrictions

File Import Restrictions

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.
Specify a single "*" to disallow the access.

Security Issue: Hosted data import of type "file" will access the files entered by the user in the context of the **server**. Therefore, in order to protect sensitive or other non-public information, you need to designate specific files and folders that users will be able to access. See the help page for more information.

Prefix-Strings of files to which access is allowed:

Default: No file access allowed

Database Import Restrictions

Define if users/groups are allowed to import hosted data from an external database:

Users/groups are allowed to import hosted data from an external database



Important: In the User Interface, the option to import the data from a text file is only available in the Member/Subscriber Importer wizard if there is at least one such prefix defined. Similarly, the option to define an importer that reads from an external database is only available if the **Users/groups are allowed to import hosted data from an external database** option is checked.

A Recipient Importer can either pull data from a database or read it from a file. In the latter case, the file can either be provided by the user at the moment the Recipient Importer is launched, or it can be read from a pre-defined location on the server (defined during importer definition). In addition, a Recipient Importer can be launched manually (from inside User Interface) or externally (with an external trigger, via a URL). Depending on the importer type and the method that is used for launching it, the following will happen:

- The Recipient Importer reads from a database and is then launched manually. As soon as the importer is launched, it reads the current data from the external database and imports it into the hosted object.
- The Recipient Importer reads from a database and is then launched with an external trigger. As soon as the importer is launched, it reads the current data from the external database and imports it into the hosted object. (The request to the trigger URL does not contain any additional data, as the data is read from the database.)
- The Recipient Importer reads from a file that is provided by the user during launch, and then the importer is launched manually. When the importer is launched, the system first asks the user to select a file to upload. This is the file that the importer is supposed to process. Once the file is uploaded (from the user's local system), the importer processes it.
- The Recipient Importer reads from a file that is provided by the user during launch, and then the importer is launched with an external trigger. The request

to the external trigger URL must be a special POST-request that already contains the file data that the importer will process. (That is, the POST-request to the trigger URL is two things at once – it is the actual trigger that starts the importer, and it is also the upload that provides the data for the importer.)

- The Recipient Importer reads from a file at a pre-defined server location, and then the importer is launched manually. When the importer is launched, it reads the file from the pre-defined location at the server and processes it. The location was pre-defined when the importer was initially created (in the Member/Subscriber Importer wizard).
- The Recipient Importer reads from a file at a pre-defined server location, and then the importer is launched with an external trigger. When the importer is launched, it reads the file from the pre-defined location at the server and processes it. The location was pre-defined when the importer was initially created (in the Member/Subscriber Importer wizard). (The request to the trigger URL does not contain any additional data, as the data is read from the file on the server.)

5.2.8 Default DomainKeys Settings

The **Default DomainKeys Settings** define whether or not to use DomainKeys signatures to authenticate the origin of the LISTSERV Maestro email messages. See Section 21 [Authenticating Message Origin with DomainKeys Signatures](#).

Figure 5-21 Default DomainKeys Settings

DomainKeys Settings

Define if e-mails that are sent by the application shall use DomainKeys signing by default. Additionally, decide on application default level if the users are allowed to decide about DomainKeys signing (using the default value defined above).
Note: These settings are only applicable if the LISTSERV host of the account or group supports DomainKeys.

Settings for Mail Jobs

Default Setting

No, do not use DomainKeys signing

Yes, use DomainKeys signing

Mail Job Specific Settings

The user must use the setting supplied above without changes for each mail job

The user may change the setting supplied above on a per-job basis

Settings for Hosted LISTSERV Lists

Enforce DomainKeys-signing for all LISTSERV Maestro standard postings to Hosted LISTSERV Lists.

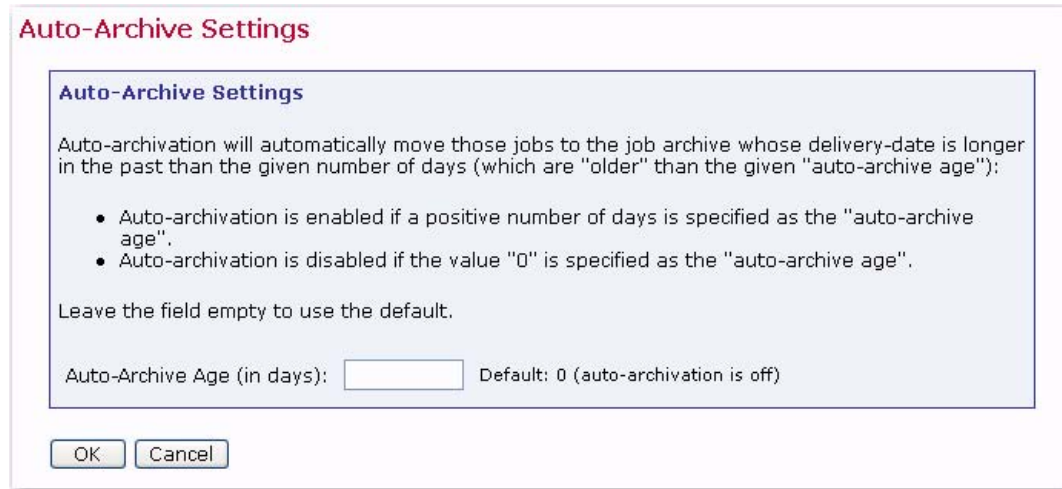
Do not enforce this and allow the user to disable or enable DomainKeys-signing on a per-list basis.

OK Cancel

5.2.9 Default Auto-Archive Settings

The **Default Auto-Archive Settings** define whether or not jobs will be automatically archived when the completed job reaches the “auto-archive age”. If the **Auto-Archive Age** is set to 0, then the auto-archive setting is off.

Figure 5-22 Default Auto-Archive Settings



5.3 Setting the Default Tracking URL

Before a user can send mail with open-up or click-through tracking, the administrator first has to configure the tracking URL that LISTSERV Maestro is supposed to use for the mail. The global default setting is used for all accounts that do not have single user or group settings. The default settings, single user settings, and group settings can be combined to offer separate tracking URL settings for different accounts and groups so that the administrator has greater flexibility in terms of customizing the tracking URL for certain accounts or groups.

When LISTSERV Maestro tracks open-up or click-through events, it does so by inserting special tracking URLs into the messages that are delivered. A typical click-through tracking URL looks something like the example below:

```
http://hostname/trk/click?ref=z4bx39x&
```

In this URL, the `hostname` points to the server where the Maestro Tracker component is installed. If the Maestro Tracker component is configured to use a non-standard HTTP port, then the tracking URL has to include the HTTP port, like the example below:

```
http://hostname:port/trk/click?ref=z4bx39x&
```

And, if the Maestro Tracker component is secured with SSL, then the HTTPS protocol would have to be used in the URL, like this:

```
https://hostname:port/trk/click?ref=z4bx39x&
```

The administrator can define the tracking URL on individual account level (for accounts that are not part of a group) or on group level (in which case, the settings are shared by all accounts in the group). Alternatively, the administrator can define default settings which are then used for all accounts/groups that do not have individual settings.

To define the tracking URL:

- **For an individual account which is not part of a group**, log into LISTSERV Maestro and go to the Administration Hub. Go to the Accounts and Identities page, and then click on the user name of the account you want to configure. Click on the **User Account** menu and select **URL Settings**. The URL Settings screen opens. Edit the URL in the **Tracking URL** section. The settings you define here will affect only the chosen account.
- **For all accounts in a group, log into LISTSERV Maestro**, log into LISTSERV Maestro and go to the Administration Hub. Go to the Accounts and Identities page, and then click on the name of the group you want to configure. Click on the **Group** menu and select **URL Settings**. The URL Settings screen opens. Edit the URL in the **Tracking URL** section. The settings you define here will affect all accounts in the chosen group.
- **To set the default for all accounts/groups that do not have individual settings** (as described in the two items above), log into LISTSERV Maestro and go to the Administration Hub. Click on the **Global Settings** menu, select **Maestro User Interface**, and then **Default URL Settings**. The Default URL Settings screen opens. Edit the URL in the **Default Tracking URL** section. The settings you define here will affect all accounts that do not have a tracking URL configured on individual account or group level.



Note: It is not possible to define the tracking URL settings on individual account level for accounts that are members of a group, e.g. all accounts in a group always share the same tracking URL settings.

Each tracking URL consists of three parts:

- **Protocol** – Usually the protocol is “http://” but can also be “https://”.
- **Host Name** – Specify the host name that is to be used in the tracking URL.
- **Port** – Specify the HTTP (or HTTPS) port that is to be used in the tracking URL, separated from the host name by a colon “:”. This part is optional if you want to use the default port for the given protocol (“80” for HTTP or “443” for HTTPS).

The full tracking URL is assembled by concatenating these three parts.

Example 1

HTTP protocol, with host “trkhost.mydomain.com” and port 8080:

```
http://trkhost.mydomain.com:8080
```

Example 2

HTTPS protocol with host “securetrk.mydomain.com” and the default port (i.e. port not specified):

```
https://securetrk.mydomain.com
```

On the user/group level, leave the **Tracking URL** setting empty to use the given default (if any). If left empty with no default given, users of this account/group will not be able to send mails with tracking.

Why does LISTSERV Maestro offer separate tracking URL settings for different accounts and groups anyway?

The reason is, to be able to offer the administrator greater flexibility in terms of customizing the tracking URL for certain accounts or groups.

One use would be if the host name of the tracker host that is used in the internal network is different than the host name that external users would have to use.

For example, if LISTSERV Maestro is installed behind a proxy, where URLs for external access need to use the proxy's host name and port, which then transparently forwards the requests to the actual tracker host behind it. In this case, the admin would specify a tracking URL using the proxy's host name and the proxy-port that is forwarded to the HTTP port on the tracker host as the **Default Tracking URL** setting that is to be used by all accounts (also see Section 24.3 [Server Name Aliases and Proxies](#)).

Customization is also useful in an environment, where the same LISTSERV Maestro setup is shared between several distinct groups, which want to discern themselves even in the tracking URL that is merged into the mails they send.

For example, the fictitious organization "MyCorp" has two divisions, one for consumer electronics and one for home appliances. The administrator of MyCorp sets up a single LISTSERV Maestro and creates two groups "electro" and "homeapp". In these groups, the admin creates accounts to be used by the members of the two divisions (which can also employ team collaboration inside of each group). The name of the server where the Maestro Tracker component is installed is "tracker.mycorp.com", but the users do not want this host name to appear in the tracking URL in the mails that are sent. Therefore, the admin also sets up two DNS names "electro.mycorp.com" and "homeapp.mycorp.com" as aliases for the same server. Finally, the administrator uses the procedure described above to set the individual tracking URLs of the "electro" and "homeapp" groups to use the corresponding host name aliases. Thus, even though both divisions are sharing the same LISTSERV Maestro installation (via their two groups), the tracking URLs generated for the mails they send are "customized" to contain a host name that fits well with each division's name.

5.4 Automatically Logging into the Maestro User Interface

The ability to automatically login to the Maestro User Interface has been added to LISTSERV Maestro 3.2 as an advanced feature and should only be configured by the system administrator. This feature can be set up for both non-group and group users.

For a non-group user:

```
http://LUIHOST/loi/index.jsp?user=USER&password=PASSWORD
```

For a group user:

```
http://LUIHOST/loi/index.jsp?user=USER&password=PASSWORD&group=GROUP
```

where *LUIHOST* is replaced with the location of your LISTSERV Maestro User Interface, *USER* is replaced with the URL-encoded version of the user name, *PASSWORD* is replaced with the URL-encoded version of the password, and *GROUP* is replaced with the URL-encoded version of the group name.

For example, assume that:

LUIHOST= example.lsoft.com

USER=holly

PASSWORD=example

GROUP=test

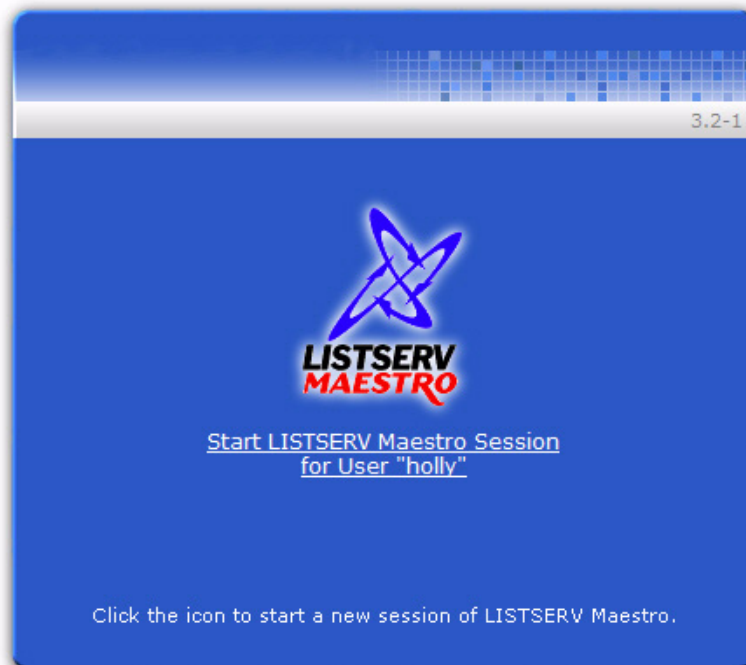
then, the following URL would be used:

<http://example.lsoft.com/lui/index.jsp?user=holly&password=example&group=test>

Once you access this URL, the Start screen opens with the **Start LISTSERV Maestro for User** link.

Click the link to start the session. The Login screen is briefly shown, and then the LISTSERV Maestro User Interface is opened and the user is automatically logged in.

Figure 5-23 Automatic Login Link



Tip: The original Start screen continues to be displayed in the background; therefore, if you click on **Start LISTSERV Maestro for User** link again, then another LISTSERV Maestro User Interface session window will be opened, with the same user logged in automatically. To prevent this from happening, you can add an additional URL parameter, "returnURL=YOURURL", where *YOURURL* is replaced with the URL-encoded version of a page-URL from your own website. The effect of this additional parameter is that once the above URL is accessed (with the "returnURL=YOURURL" parameter included), then the Start screen will be displayed; however, when you click the **Start** link, two things will happen: The LISTSERV Maestro User Interface session will be opened, with the user automatically logged in (just as above). In addition, the original Start screen will be returned back to the *YOURURL* given in the "returnURL" parameter.

Section 6 Administrative Policies

Every institution and business using networks computers will have its own administrative policies for data backup, error reporting, software access on the network, user accounts, and so on. Assimilating LISTSERV Maestro into the existing administrative structure is done through **General Administration** and **Administrative Email Notification** settings.

General Administration settings are available for each of the three components, HUB, LUI, and TRK. Each component can have its own settings for saving backups and logging activity. The Administration Hub component has additional general administrative settings for external processes to be run after a backup completes (see Section 11.2 [Configuring External Post-Backup Processes](#) for details). The Administration Hub also has a link to enable special email notification messages to be sent to a configured address or addresses in the event of a system problem and/or a system startup (see Section 6.4 [Administrative Email Notifications](#) for details). The Maestro User Interface has additional settings for transferring events from the Maestro Tracker component, job archiving, and runtime administration.

6.1 Configuring Backups

The Administration Hub component of LISTSERV Maestro acts as the backup master for all three components. This way, any problems that might arise from having different components that store data independently and reside on different servers is avoided.

The Administration Hub will centrally trigger a backup on all connected components (including itself) in order that the backup data saved by each component is consistent with the backup data of all other components. This backup is initiated based on the values entered in the Global Component settings for the Administration Hub. See Section 11 [Saving and Restoring a Backup](#) for additional information.

Each component can save its backup in its own folder configured in the **General Component Settings** section of the component. See Section 11.3 [Configuring the Backup Location](#) for more information.

6.2 Runtime Administration and System Shutdown

These settings allow the administrator to influence the availability of the Maestro User Interface. For example, they can be used in preparation of the system shutdown for maintenance by disallowing new logins, disabling the Outbox to prevent outgoing jobs from being sent, and sending a warning message to users already logged on to the system.

Runtime Administration settings are found by clicking on the **Global Settings** menu, selecting **Maestro User Interface**, and then finally selecting **General Administration**. Use these settings to reach a safe shutdown state, where shutting the system down will not disrupt jobs delivery or users in the midst of preparing a job. Follow these steps to safely shutdown LISTSERV Maestro:

1. Open a web browser and access the Maestro User Interface, logging in with the administrative account. See Section 8 [Special Administrative User Account](#) for more information.
2. Open a second browser and access the Administration Hub. Go to **Global Settings > Maestro User Interface > General Administration**. The General Administration of Maestro User Interface screen opens.
3. Check the **Sending is disabled** option. This will stop any new jobs from starting their send process.
4. Check the **LISTSERV Maestro User Interface is locked** option. This will stop any new users from logging in.
5. In the top text box, enter an informational message such as "The system will shortly go down for maintenance, therefore login is currently not possible" so users denied login will know why.
6. In the bottom text box, enter an informational message such as "The system will shortly go down for maintenance. Please finish your current work and log out as soon as possible or contact the administrator" so that all current users will now see this warning at the top of every page they access and will know to wrap up whatever they are doing. Consider adding a message telling users exactly when the system will shutdown and for how long.
7. In the browser that is logged into the administrative user account, go to the Ongoing Jobs screen. All of the pending jobs will be listed in a table. Click on the **State** link to sort the jobs according to their processing status. Jobs that are in the process of sending will be indicated by a yellow arrows icon. Refresh the screen to renew the list. When the jobs have finished processing, they will no longer appear in the table after refreshing the screen. New deliveries will not start since the **Sending is disabled** option in the Administration Hub has been set.
8. After all jobs that were in the process of sending have finished and currently logged in users have had enough time to wrap up what they were doing, LISTSERV Maestro can be safely shut down and maintenance tasks can be executed.
9. After the restart, return to the HUB and uncheck the **Sending is disabled** and **LISTSERV Maestro User Interface is locked** options to make the Maestro User Interface available and working normally again.



Notes: When the LISTSERV Maestro Tracker component is shut down, all tracking URLs becomes unavailable, and all other tracking activity stops. Mail recipients will not be able to click on links in the message and no tracking events will be recorded. If at all possible, install the tracker component on its own server to minimize down time. Try to schedule system shutdowns at a time that disrupts the fewest users and the fewest possible mail recipients.

LISTSERV Maestro can be put into "Maintenance Mode" by setting the INI file `MaintenanceMode = true` and then restarting the component. While in this mode, no users will be able to log into the Maestro User Interface, so maintenance tasks can be executed. To return the component to normal, change the setting to `false`, delete the entry, or comment it out with a leading `"#"` or `"!"` and restart.

6.3 User Restrictions

LISTSERV Maestro has many features that allow regular users' activities within the system to be limited. Some limitations occur on a system-wide level, such as not allowing multiple logins from the same account and some limitations can be configured to occur on a system, group, or individual level.

User access to LISTSERV Maestro can be limited to a single login per account or allow multiple logins per account. This setting is located in the **Runtime Administration** section of the General Administration of Maestro User Interface screen. See Section 17.2 [Disallowing Concurrent Access with the Same User Account](#) for more information.

In addition, an advanced security option lets you limit the number of invalid login attempts, and when this is surpassed, lock the account and deny access. LISTSERV Maestro supports this form of login locking in the Administration Hub and in the User Interface component. For more information, see Section 17.3 [Securing Access Against Dictionary Attacks](#).

Each of the LISTSERV Maestro components (HUB, LUI, and TRK) can be configured to restrict access based upon the IP address of the computer where the browser/email-client is running that is used to access the component. This means that it is possible, for example, to define that everyone (all IP addresses) is allowed to access the Maestro Tracker component, but only certain addresses (a local subnet, perhaps) are allowed to access the Maestro User Interface and Administration Hub components. See Section 17.1 [IP Address Restrictions](#) for more information.

Other user restrictions that can be configured at the system level are:

- **Maximum size limit for an email message** – Sets a limit for the total byte size of the message after all transfer encoding and MIME multipart wrappers have been applied. This setting can be set as the default in the **Global Component Settings > Maestro User Interface > Default Size Limits**. The default can be overridden by setting this limit at the group or user level when administering user accounts. See Section 7.4 [Editing Account Information and Assigning Single User Settings](#).
- **Maximum file size for uploaded files** – Applies to all types of files uploaded to the system including recipient lists, HTML and text messages, attachments, images, and so on. This setting is only available application wide. It is not available at the group or user levels.
- **Content Restrictions** – Allows special AOL Rich Text formatting as part of defining an email message. If allowed, users can choose to include an AOL alternative in any HTML message created. This setting is obsolete and not recommended.
- **Drop-in Content Restrictions** – Creates a positive list for files and a positive list for URLs that are going to be used as drop-in content elements. This helps prevent security breaches into local files and URLs. If this setting is left blank on the system level, it must be set on the group or user level to allow those accounts to use files and/or URLs as drop-in content. If left blank on every level, drop-in content of these types will not be allowed. See Section 7.4 [Editing Account Information and Assigning Single User Settings](#) for more information.
- **Recipients Restrictions** – The top section, **Recipients Type Restrictions**, sets the type of recipients the user is able to use for an email job. If **Disabled** is selected, the option will appear grayed out in the Maestro User Interface, and the user will not

be able to select it. If **Hidden** is selected, the recipient type will be disabled and will not appear at all in the Maestro User Interface. The default can be overridden when configured at the group and user level. See Section 7.4 [Editing Account Information and Assigning Single User Settings](#) for more information.

The bottom section, **Recipients Upload Restrictions** contains a text box for the administrator to enter allowable paths or path prefixes leading to files on a server accessible to the Maestro User Interface. These files are used for uploading "just before sending" CSV files for recipients definitions. If left blank, using these CSV files will not be allowed.

- **Hosted List Restrictions** – Under **Global Settings > Maestro User Interface > Hosted Data Settings**, check or uncheck boxes to allow users and groups to create Hosted Recipient Lists and Hosted LISTSERV lists. If the boxes are left unchecked, users and groups may not create such lists. If Hosted LISTSERV Lists are allowed, their connection settings need to be configured under the Connection Settings for Hosted LISTSERV Lists (see Section 4 [LISTSERV and LISTSERV Maestro](#)).
- **Tracking Restrictions** – Defines the type of tracking that will be available, by default. This setting can be set as the default under **Global Settings > Maestro User Interface > Default Tracking Restrictions**. To enable a tracking type, select **Enable**. To disable, but still display the tracking type, select **Disabled**. To disable and hide a tracking type, select **Hidden**.

Other restrictions can be placed on individual accounts when configuring **Team Collaboration** settings. Team collaboration settings allow the job owner to give or revoke privileges to group members affecting their abilities to create jobs, work on particular parts of jobs like defining recipients, and use jobs in reports. These settings can be configured at a default level for all jobs that an account owns under user settings for an account, and they can be set at the job level for individual jobs.

6.4 Administrative Email Notifications

LISTSERV Maestro can send email messages to one or more email addresses in the event of a system problem or system startup. Once configured, errors and/or startups that occur on any component will trigger a message. If an error occurs on three components, three separate messages will sent to each configured recipient address.

In addition, the option to send an email after system problems also controls whether or not to send notification after each backup to inform you of its success or failure.

Finally, LISTSERV Maestro also gives you the ability to define different SMTP servers and ports used for the email notifications for each component.

To configure administrative email messages, click on the **Global Settings** menu, select **Administration HUB**, and then finally select **Administrative E-mail Notifications**.

To have administrative email notifications sent, select the option to send email notifications. If notification is desired for system start, select **Send a notification email in the event of sever system problems, and optionally:**, and then check **Send notification e-mail at each startup**. The following settings need to be configured to use email notifications:

- **SMTP Host** – Enter the host name running the SMTP server that will be used for the mailing. This field is mandatory and must be filled out with a valid host name that

can be used to send e-mail notification by the LISTSERV Maestro component that encounters a severe problem.

- **SMTP Port** – Enter the SMTP port on which the SMTP server on the host specified above listens for SMTP connections. This field is optional; if left empty, the standard SMTP port "25" is used.



Important: If you only supply the default SMTP host and port in this section, then this host / port must be reachable from all servers running one of the LISTSERV Maestro components. If not all component servers can reach this host by using this name and port (for example due to firewalls or DNS or other networking-related issues), then supply the custom SMTP Host / Port (see below for details).

In the **Custom SMTP Hosts and Ports for LUI / TRK** section on the Administrative Email Notifications screen:

- To use the default SMTP host name and port supplied in the section above, then select the **Use Default SMTP host and port as defined above** option from both the **LUI** and **TRK** drop-down menus.



Note: The Administration HUB always uses the default SMTP host name and port defined in the section above.

- If the default SMTP host name and port supplied in the section above can not be used for all LISTSERV Maestro components, then use the settings in this section to supply SMTP host names and/or ports for LUI and/or TRK.
 - **LUI SMTP Host / Port** – If you want to supply a custom SMTP host and/or port for the Maestro User Interface component, then choose the **Use SMTP Host and port as supplied below** option, and then supply SMTP host/or port that the LUI component will use for notifications.
 - **TRK SMTP Host / Port** – If you want to supply a custom SMTP host and/or port for the Maestro Tracker component, then choose the **Use SMTP Host and port as supplied below** option, and supply the SMTP host/port that the TRK component shall use for notifications.
- **Sender Address** – Enter a sender address that will be used as the sender address for all the email notifications. This field is mandatory and must be filled out with a valid Internet email address.
- **Notification email will be sent to the following addresses** – Specify at least one valid Internet email address that will be the recipient of the notifications sent from LISTSERV Maestro. This field is mandatory. Multiple addresses can be entered, one per row, with no separator characters. All addresses entered here will appear in the "To:" field of the email notification, so each recipient will be able to see the addresses of all other recipients.



Tip: You can create a LISTSERV List to distribute notifications as well as archive them. Simply use the list address in this field and add as many addresses as needed to the LISTSERV List.

Figure 6-1 Administrative Email Notifications

Administrative E-mail Notifications

Administrative E-mail Notifications

Do not send administrative e-mail notifications

Send e-mail notifications in the event of severe system problems, and optionally:

Send a notification e-mail at each startup.

The following settings will be used to send the notification e-mail:

SMTP Host / Port

SMTP Host: <default undefined>

SMTP Port: Default: 25

Custom SMTP Hosts and Ports for LUI / TRK

LUI:

TRK:

E-Mail Addresses

Sender Address: <default undefined>

Notification e-mail will be sent to the following addresses:
(Specify one per row)

<default undefined>

Send a test e-mail to the addresses listed above when this page is submitted with "OK"

OK Cancel

6.4.1 Testing Email Notifications

It is important to test the settings for email notifications to make sure that they do work, and that the specified addresses receive the mail sent by the system. This verification is done with the **Send a test email to the addresses listed above** checkbox.

Checking this option, and then submitting the page by clicking **[OK]**, will send a test email to all recipients specified. A test email will be generated by each of the LISTSERV Maestro components so that each of the addresses will receive three different test messages, one from each component.

As the next step, verify that all specified addresses received three test-notification email messages. If this is not the case, then the notification sending needs some troubleshooting. Follow these steps to troubleshoot the email notification settings:

- Check the log file(s) of the component(s) that did not send email notification. Verify that the log(s) contains an entry with the following text:

- “Administrative email notifications have been enabled.
This message is for testing that administrative email notifications have been enabled correctly, there is no problem with the application!”
- If this message does not appear, then the **Send a test email to the addresses listed above** checkbox was not actually checked when the page was submitted, or the **[OK]** button was not clicked and the screen was exited by the **[Cancel]** button or any of the shortcut icons.
- If the message above appears in the log file, then check the log file for an error message that appears right after the message quoted above. The error message will read: “*Error when trying to send notification email about previous log entry: Error description here...*”
- The error description will provide an idea of what needs to be changed to make the messaging work (for example, the error could be caused by an incorrect host name or SMTP port).
- If the first message appears in the logs, but not the second (the error message), LISTSERV Maestro presumes the email notification was successfully sent. If this happens, take a closer look at the SMTP server and the other components in the mail delivery chain to find out where the mail got lost.

6.5 Refreshing the Subscriber Page Translations

Translations for the Subscriber Pages are maintained externally with the L-Soft Resource Translation Tool and are made available to LISTSERV Maestro as `JAR` files. New versions of these files are loaded by the system during startup. If you want to reload the translation `JAR` files without restarting the system, open a browser and access the Administration Hub. Then, click on **Global Settings > Maestro User Interface > General Administration**. The General Administration of Maestro User Interface screen opens. In the **Subscriber Access Page Translations** section, click the **[Refresh Translations Now]** button.

For more information on the Resource Translation Tool, see the Interface Customization Manual for LISTSERV Maestro 4.0.

Section 7 Creating and Editing User Accounts and Identities

A user account is used for the actual login to LISTSERV Maestro. To be able to login, a user must supply a valid account name (and the account's group name, unless the account does not belong to a group), together with the account's password. Each user account is associated with a certain set of user rights that allow or disallow different actions in LISTSERV Maestro while you are logged in with this account.

An identity is a collection of several accounts which belong to one and the same "identity", usually a person. By collecting all accounts of one person into an identity, LISTSERV Maestro knows that these accounts all belong together. As a result, the user is then allowed to switch between the accounts in the identity without having to perform an actual logout and login. In other words, if a user logs in with one account that belongs to an identity, he can then switch over to all other accounts in the same identity without having to first logout the old account and then login again with the new account (he does not even have to provide the password of the new account). Identities are useful in cases where one and the same person was assigned several accounts for different purposes. The user can then simply log in with one account and perform tasks under this account, then easily switch to any of his other accounts.



Important: Because of this switching between accounts in the identity, it is important not to add any accounts to one and the same identity, between which such a switching will not be allowed.



The Accounts and Identities screen displays a list of all defined accounts and identities.

Figure 7-1 Administering User Accounts and Identities

Group	Account	Identity
quest	bparker	ben
quest	francoise	FMB
quest	jharlan	john
quest	liam	liam
quest	mshannon	Michael

Identity	Accounts
FMB	francoise, guest/francoise, sample/francoise
HT	holly, sample/holly
Michael	mshannon, guest/mshannon, sample/mshannon
Noodles	noodles, sample/harpo, sample/noodles, test/noodles
ben	bparker, guest/bparker, sample/bparker
john	jharlan, guest/jharlan, sample/jharlan
liam	liam, guest/liam, sample/liam

Each user name is a link that opens the Define User Account screen, and each identity name is a link that opens the Define Identity screen. Links on the Define User Account screen lead to user settings that can be changed for only that single user account. If the user is a member of a group, the group name is a link that opens the Group Settings screen, containing settings that can be changed for the entire group.

From the Accounts and Identities screen it is possible to:

- Add a new user account or identity
- Assign users to a group or identity
- Edit existing user account and group settings, including deleting the accounts, identities, or groups

7.1 Creating a New User Account



To create a new user account, click on the **Accounts and Identities** menu and select **New Account**. The Define User Account screen opens.

Figure 7-2 Defining User Account

- **Name** – Fill in a user name. User names are not case sensitive, and can be composed of letters, numbers and symbols. User names can also contain spaces. User names must be unique if users are not assigned to a group. Users assigned to different groups may have the same user name as long as the combination of user name and group name is unique.
- **Group** – Assign the user to a group, if desired. Groups are optional and can be left blank. However, users must be assigned to a group in order to use the team collaboration features. Only members of the same group can collaborate on email jobs. The combination of user name and group name must be unique. There cannot be two users with the same name in the same group, although there may be two users with the same name in different groups.
- **Identity** – Assign the user to an identity, if desired. Identities are optional and can be left blank.

- **Password** – Assign a password. All passwords are case sensitive, and must be at least five characters long.
- **Confirm Password** – Retype the password to confirm it.

Check **The user is allowed to change his password** if you want to grant the user permission to change his or her password. Uncheck this option to revoke this privilege. Click **[Save]** to save the account information and continue. Click **[Cancel]** and the new user will not be created.



Note: If the user is allowed to change his/her password, then the **Change Password** option on the **Utility** menu in the LISTSERVE Maestro User Interface will be available.

7.2 Creating a New Identity



To create a new identity, click on the **Accounts and Identities** menu and select **New Identity**. The Define Identity screen opens.

Figure 7-3 Defining Identity

Enter the identity's name in the **Name** field, and then select the accounts that are in the identity. To add an account to the identity, select it in the **Available Accounts** list and click the **[->]** button. This will move the identity into the **Accounts in Identity** list.

To remove an account from the identity, select it in the **Accounts in Identity** list and click the **[<-]** button. This will move the identity into the **Available Accounts** list.

Click **[OK]** to submit your changes, or click **[Cancel]** to exit without submitting your changes.



Important: LISTSERV Maestro interprets all accounts in an identity as related and allows a user to switch between these accounts without providing the login information for each account. In other words, if a user logs in with one of the accounts in the identity, then this user can switch to all other accounts in the same identity without having to know the passwords of these other accounts. Because of this, it is very important that you do not combine accounts into one identity that should not be accessed by the same user.

7.3 Managing User Rights



To define the Maestro User Interface user rights, click on the **Global Settings** menu and select **User Rights**. The User Rights Management for Maestro User Interface screen opens.

Figure 7-4 User Rights Management Screen

User Rights Management for Maestro User Interface
Use the checkboxes to define the user rights of the corresponding accounts.

Displayed accounts:

User	Create Jobs	Create Reports	Admin Sender Profiles	Admin Drop-Ins	Admin Content Templates	Admin Target Groups	Admin Datasets	Link Datasets & WA	Designated Job Owner
bparker	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<This Account>
francoise	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	jharlan
jharlan	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<This Account>
liam	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<This Account>
mshannon	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	jharlan

From this screen, you can assign user rights for an entire group or for a specific user in the selected group.



Tip: You can also click on a specific account in the Accounts and Identities screen and change the Maestro User Interface rights for the selected account only. For more information, see Section 7.4 [Editing Account Information and Assigning Single User Settings](#).

The following user rights can be assigned or unassigned:

- **Create Jobs** – Grants the right to create new jobs. If the user is a member of a group with the right to create new jobs, it is necessary to also define who owns the jobs that are created by this account. Jobs can be owned by the account that creates them or by another group member selected from the **Designated Job Owner** drop-down menu. If jobs created by one account are owned by another group member, when the user creates a new job, the ownership will be set to the other group member and the team collaboration default preferences of that account are applied.



Important: If the owner account of a job that another account tries to create has not given that account at least one right in their team collaboration default settings, then the creating account will not be able to start a new job, and an error message will appear on the Start New Job screen.

- **Create Reports** – Grants the right to create new reports. Existing reports are available for all members in a group.
- **Admin Sender Profiles** – Grants the right to create new sender profiles. Existing sender profiles are available for all members in a group.

- **Admin Drop-Ins** – Grants the right to create new drop-in content elements. Existing drop-ins are available for all members in a group.
- **Admin Content Templates** – Grants the right to create, edit, and delete message templates. Enabled templates are available for all members in a group.
- **Admin Target Groups** – Grants the right to administer existing target groups and to create new recipient target groups by providing access to Recipients Target Groups Wizard.
- **Admin Datasets** – Grants the right to administer recipient datasets in the recipient warehouse. This includes creating, editing, and deleting datasets, hosted lists, and individual subscribers. It also grants the right to administer target groups.
- **Link Datasets & WA** – If granted, the user may create a link between any recipient dataset and the LISTSERV Web Interface of the LISTSERV instance that the user is connected to. For a dataset that is linked, the membership area of the dataset will automatically provide links that allow the subscribers to directly access the list archive pages of the LISTSERV Web Interface.



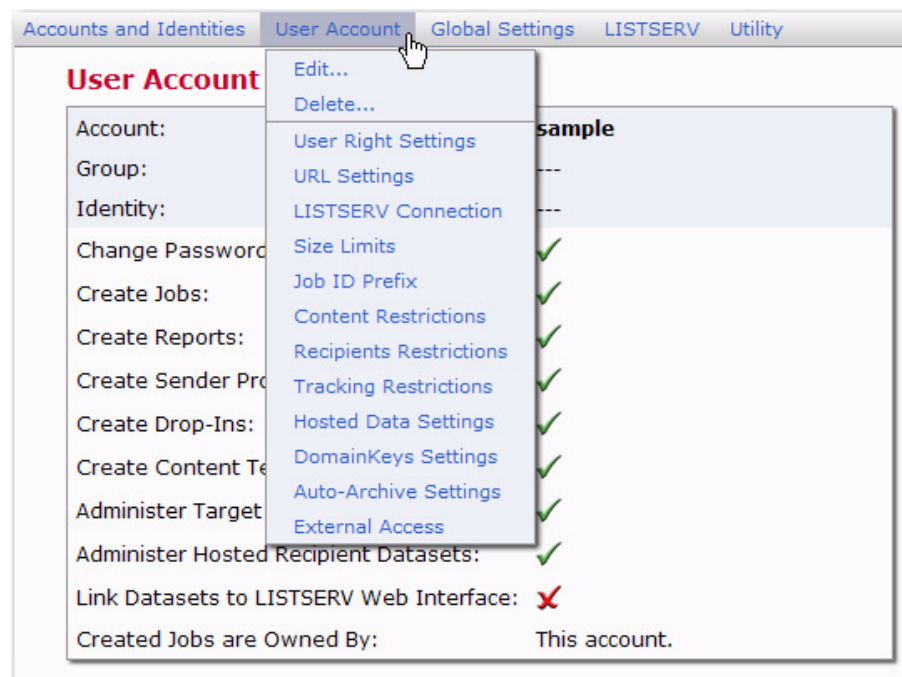
Tip: By clicking on the **User** link, then every member of the group will be granted every user right. In addition, by clicking on the name of the user right (e.g. **Create Jobs**), then every member of the group will be granted that specific user right.

7.4 Editing Account Information and Assigning Single User Settings



To edit existing account information, go to the Accounts and Identities screen (by clicking on **Accounts and Identities > Account and Identities Overview**). Click on the account name you want to edit. The User Account Overview screen opens. From here, the **User Account** menu is available, giving you access to every component of the user's account.

Figure 7-5 Editing User Account Settings



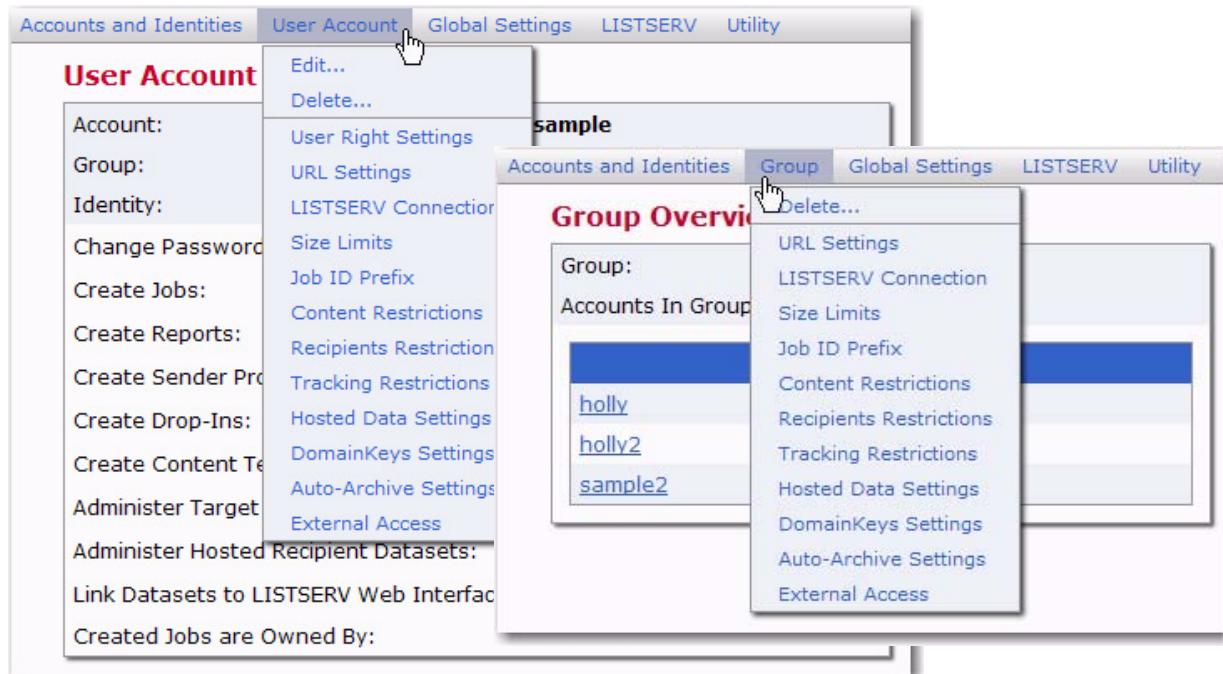
7.4.1 Editing General User Settings

New To change the user name, group, identity, or password of an existing account, click on the **User Account** menu and select **Edit**. The Define User Account screen opens. To keep an old password while changing other settings, leave both password fields blank.

7.4.2 Editing Component Specific Settings for Single and Group Users

New The **User Account** and **Group** menus make it easy to administer specific settings for the the selected account or group.

Figure 7-6 User Account and Group Menus



Click on any of the menu(s) options to open a screen to edit that particular setting. The options and settings on each menu vary slightly. The main difference is that settings configured at the group level will affect all members of the group; whereas, settings configured for the single user will only affect that user.

Settings can include any of the following:

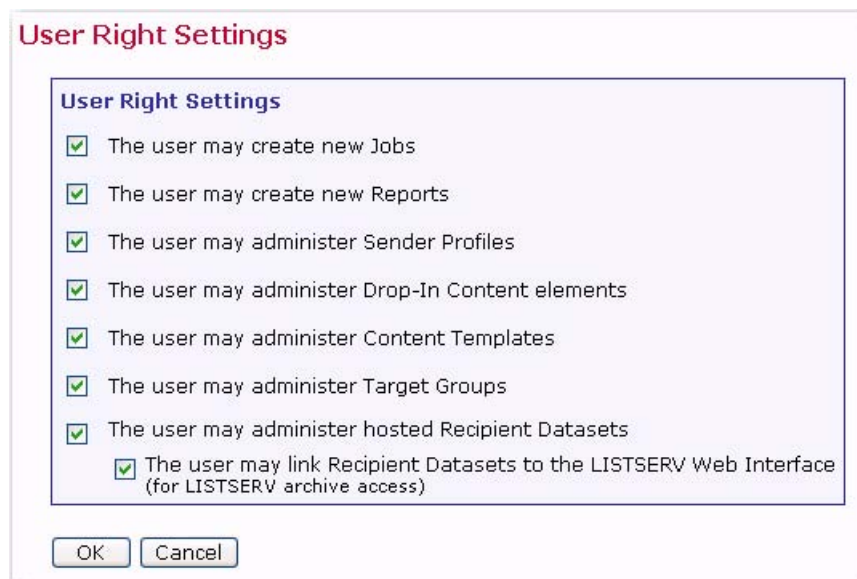
- **User Right Settings** – The user rights apply only to the configured user, even if the user belongs to a group. Check the boxes next to the privileges to be granted to the user. Uncheck the boxes next to the privileges to be revoked from the user. User rights settings include:
 - **The user may create new Jobs** – Grants the right to create new jobs. If the user is a member of a group with the right to create new jobs, it is necessary to also define who owns the jobs that are created by this account. Jobs can be owned by the account that creates them, or by another group member, selected from a drop-down menu of existing accounts in the group. If jobs created by one account are owned by another group member, when the user creates a new job, the ownership will be set to the other group member and the team collaboration default preferences of that account are applied.



Important: If the owner account of a job that another account tries to create has not given that account at least one right in their team collaboration default settings, then the creating account will not be able to start a new job, and an error message will appear on the Start New Job screen.

- **The user may create new Reports** – Grants the right to create new reports. Existing reports are available for all members in a group.
- **The user may create new Sender Profiles** – Grants the right to create new sender profiles. Existing sender profiles are available for all group members.
- **The user may create new Drop-In Content Elements** – Grants the right to create new drop-in content elements. Existing drop-ins are available for all members in a group.
- **The user may administer Content Templates** – Grants the right to create, edit, and delete message templates. Enabled templates are available for all members in a group.
- **The user may administer Target Groups** – Grants the right to administer existing target groups and to create new recipient target groups by providing access to Recipients Target Groups Wizard.
- **The user may administer hosted Recipient Datasets** – Grants the right to administer recipient datasets in the recipient warehouse. This includes creating, editing, and deleting datasets, hosted lists, and individual subscribers. It also grants the right to administer target groups.
 - **user may link Recipient Datasets to the LISTSERV Web Interface** – If granted, the user may create a link between any recipient dataset and the LISTSERV Web Interface of the LISTSERV instance that the user is connected to. For a dataset that is linked, the membership area of the dataset will automatically provide links that allow the subscribers to directly access the list archive pages of the LISTSERV Web Interface.

Figure 7-7 User Right Settings



- **LISTSERV Connection** – The LISYSERV Connection can be set at several levels. The widest level of setting is the Global Component Setting, which defines the global default LISYSERV Connection on the application level (see Section 5.2 [Application Default Settings](#)). These settings will be used as defaults for all accounts that do not have individual settings for group or user level.

The next level is the default setting for a group. To set the LISYSERV settings for an entire group in the user list, click on the group name (appears only if the account belongs to a group). If defined on the group level, the group settings will override the global default settings for that group. The settings will apply to all accounts in the group, except for the **LISTSERV Client Address** and **LISTSERV Client Password** settings. These two settings may be configured individually for each account in the group. Accounts in the group for which they are not configured will use the group's settings as the default.

Individual accounts are the last level of settings for the LISYSERV Connection. Individual accounts can belong to a group (group user) or not belong to a group (single user). Settings for individual accounts will override default group and global settings. To set the LISYSERV Connection for an individual account, click on the user name from the user list. The screen that opens is different depending on whether the account is a group user or a single user. Single user accounts can have all LISYSERV Connection settings defined. Group user accounts can only define the **LISTSERV Client Address** and **LISTSERV Client Password** settings. These individual group settings result in a different LISYSERV email address for each user so that jobs can be identified by owner in the LISYSERV logs.

For information and instructions on how to fill out the fields for setting the LISYSERV Connection, see Section 4 [LISTSERV and LISYSERV Maestro](#). For information about setting a special external host name, see Section 4.2.1 [Specifying the LISYSERV Host with Different Internal and External Names](#). For information on setting up a dedicated LISYSERV instance for processing bounces, see Section 4.2.2 [Specifying a Separate LISYSERV Instance for Processing Bounces](#).



Notes: To prevent inadvertent sharing between groups, each group should have a separate LISYSERV connection client address and password. For extra security, use a separate LISYSERV instance for each group. Remember that each LISYSERV instance requires a separate license, but multiple LISYSERV instances can run on a single Windows server.



- **URL Settings** – This defines the settings for the various URLs used by LUI.

Note: If defined on the application level, then these settings will be used as defaults for all accounts that do not have individual settings on the group or user level. For details, see Section 5.2.1 [Default URL Settings](#).

Figure 7-8 URL Settings

URL Settings

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

URLs for User

LUI Access URL for Users
 This LUI access URL is used by normal LUI users to access the LUI user interface.
 Important: A user must specify exactly the protocol, hostname and port defined here to be allowed to login to LUI.

Default: http://<Host name of server running the LUI component>

LUI Access URL Aliases for Users
 These URLs are aliases for the LUI access URL.
 If the user accesses any of the alias urls, he is redirected to the official URL defined above.

Default: <none>

Subscribers Access URL
 The subscribers access URL is used by subscribers to access the subscriber pages (membership area pages).
 Important: A subscriber must specify exactly the protocol, hostname and port defined here to be allowed to login.
 For this URL, the same URL is used as for the LUI access URL defined above, unless it is overridden here:

Default: http://<Host name of server running the LUI component>

Subscribers Access URL Aliases
 These URLs are aliases for the subscriber access URL.
 If the user accesses any of the alias urls, he is redirected to the official URL defined above.

Default: <none>

HUB Access URL for Users
 The HUB access URL is needed by the system whenever a normal user changes his password.
 This URL is not used by the users directly.

Default: http://<Host name of server running the HUB component>

Tracking URL
 The tracking URL is used to generate the tracked links in an e-mail.
 Tracking is disabled if no tracking URL is specified, either directly or by default!

Default:

- **LUI Access URL** – The LUI access URL is used by the LUI users and the admin, to access the LUI user interface. For some LUI interface pages, the

system needs to include this access URL on the page itself (for example to create a link). The URL that the system uses for this is generated using the settings supplied here. If a setting is not supplied, the displayed default value will be used.

Specify the **Host Name**, **HTTP Port**, and **Protocol** (either HTTP or HTTPS) that is to be used to generate the LUI access URL.

Only on user/group level: If a LUI server has several host names assigned, then normally a user can use any of these host names to access LUI. If you want to restrict the user to only be able to login when he uses the host name as specified here, check the User must use this host name for login option. If checked, then a login by the user to which this setting applies will only be successful, if the user has used the same host name in his access URL as is specified here. If the user uses a different host name in his access URL, the system will behave as if the user account does not exist.

- **LUI Access URL Aliases for Users** – The LUI Access URL Aliases will be the aliases for the LUI Access URL. If a user accesses LUI with any of the alias urls, then their login is accepted, and after login, they will be automatically redirected to the official LUI Access URL.

Specify a list of alias URLs (one per row), each with the full protocol and hostname, optionally followed by the port. For example,

```
http://luihostalias.domain.com
```

or

```
https://secureluihostalias.domain.com:1443
```

- **Subscribers Access URL** (only available on user/group level) – The subscriber pages access URL is used by subscribers to access the subscriber pages (the pages of the membership area of a dataset). There are various situations in which LUI needs to generate this URL, for example to include an unsubscribe link, login link or forward-to-a-friend link in an e-mail message. The URL that the system uses for this is generated using the settings supplied here. By default, the system will simply use the same settings as are supplied for the LUI access URL (see above). You only need to specify separate values for the subscriber pages access URL if this URL is actually supposed to be different than the LUI access URL.

If necessary, specify the Host Name, HTTP Port and Protocol (either HTTP or HTTPS) that is to be used to generate the subscriber pages access URL.

- **Subscribers Access URL Aliases** (only available on user/group level) – The subscriber access URL aliases are the URLs that will be the aliases for the Subscriber Access URL. If the subscriber accesses any of the alias urls, then they are automatically redirected to the official Subscriber Access URL. By default, the system will simply use the same alias URLs that are supplied for the LUI access URL (see above). You only need to specify separate alias URLs for the subscriber pages if these aliases are actually supposed to be different than the LUI access aliases. Make sure to specify a list of alias URLs (one per

row), each with the full protocol and hostname, optionally followed by the port. For example,

```
http://subscriberhostalias.domain.com
```

or

```
https://securesubscriberhostalias.domain.com:1443
```

- **HUB Access URL** – The HUB access URL is used behind the scenes to allow the system to display the password change page for normal users, and to allow the admin to switch seamlessly between the LUI and the HUB interfaces. To enable this functionality, the system needs to be able to generate a URL that correctly points to the HUB user interface. The URL that the system uses for this is generated using the settings supplied here. If a settings is not supplied, the displayed default value will be used.

Specify the **Host Name**, **HTTP Port**, and **Protocol** (either HTTP or HTTPS) that is to be used to generate the HUB access URL.

- **Tracking URL** – The tracking URL is used to generate the tracked links in e-mail messages for which tracking has been enabled. The URL that the system uses for this is generated using the settings supplied here. If a settings is not supplied, the displayed default value will be used.

Specify the **Host Name**, **HTTP Port**, and **Protocol** (either HTTP or HTTPS) that is to be used to generate the tracking URL.



Important: There is no default for the tracking URL host name. Therefore, as long as no such hostname is defined, tracking is disabled because a tracking URL can not be generated.



Tip: For more details on tracking URLs, see Section 5.3 [Setting the Default Tracking URL](#).

- **Size Limits** – Sets a size limit for email messages. The size limit for a message applies to the total byte size of the message after all transfer encoding and MIME multipart wrappers have been applied. If the message exceeds the size limit, the delivery will fail.

Figure 7-9 Size Limits

- **Job ID Prefix** – A Job ID Prefix is an optional part of the Job ID. The prefix comes before the system generated Job ID and is separated by a hyphen. Allowing, disallowing, or presetting Job ID Prefixes for users takes place by selecting the

corresponding option button and entering the prefix. For more information on job ID prefixes, see the LISTSERV Maestro Users Guide.



Tip: Use preset Job ID prefixes to identify jobs from different groups or users more easily. For internal charge-back purposes, the Job ID can be extracted from the job name and base charges on records in LISTSERV's system changelog. If no such identification is necessary, select **Any job ID prefix allowed** to give users a way of grouping jobs.

Figure 7-10 Job ID Prefix Settings

- **Content Restrictions** – This option defines restrictions for the content of email messages. In the top section of the screen, select the option button to allow or disallow an AOL formatted alternative for HTML email messages. For more information on HTML messages, see the LISTSERV Maestro User's Manual. The AOL alternative is obsolete and not recommended.

Figure 7-11 Content Restrictions

In the bottom section of the screen, create a “positive list” of all files and/or URLs that will be available for drop-in content. If the list is left blank, no drop-in content in the form of files and/or URLs will be allowed. See the online help for more information on using this setting.

Figure 7-12 Drop-In Content Restrictions

Drop-In Content Restrictions

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

Security Issue: Drop-In content elements of type "file" and "URL" will access the files or URLs entered by the user in the context of the **server**. Therefore, in order to protect sensitive or other non-public information, you need to designate specific files, folders, and URLs that users will be able to access. See the help page for more information.

Prefix-Strings of files to which access is allowed:

Default: No file access allowed

Prefix-Strings of URLs to which access is allowed:

Default: No URL access allowed

- **Recipients Restrictions** – In the top section of the screen, set the type of recipients available for use for an email job. If **Disabled** is selected, the option will appear grayed out in the Maestro User Interface, and the user will not be able to select it. If **Hidden** is selected, the recipient type will be disabled and will not appear at all in the Maestro User Interface.

To limit the number of email jobs a group or user sends per day and/or month, click the **Enable** option next to any of the permissible recipient types (see list below). Next, in the **Mail Limits** fields, enter the per day and per month limits. This limit is used to prevent excessive limits for those wishing to control the amount of email being sent (e.g. for server limitation reasons or for license limitations).

Limiting the number email jobs sent can only be defined for the following recipient types:

- **Upload a Recipients Text File** (Limits can be defined for both Recipients and Target Groups.)
- **Select Recipients from a Database** (Limits can be defined for both Recipients and Target Groups.)
- **Determine Recipients Based on Reaction on Another Job** (Limits can only be defined for Recipients.)
- **Sent to a Hosted Recipient List** (Limits can only be defined for Target Groups; in addition, they are automatically enabled if a HRL already exists.)



Note: Recipient types that are controlled by LISTSERV (**Send to an Existing LISTSERV List** and **Let LISTSERV Select from a Database or LDAP Directory**) cannot have email limitations applied to them.

Figure 7-13 Recipients Restrictions - Type Restrictions

Recipients Restrictions

Recipients Type Restrictions

Define which recipients types are available by default.
 Select "Enabled" to enable a recipients type.
 Select "Disabled" to disable but still display a certain recipients type.
 Select "Hidden" to disable and hide a certain recipients type.

Recipients types managed by LISTSERV Maestro	Enabled	Disabled	Hidden
Upload a Recipients Text File (Recipients and Target Groups)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Select Recipients from a Database (Recipients and Target Groups)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Determine Recipients Based on Reaction on Another Job (Recipients only)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Send to a Hosted Recipient List (Target Groups only)	(automatically enabled if a HRL is present)		

Supply the mail limits for the recipients types above. Leave the fields empty to inherit the default values, supply 0 to use unlimited mail delivery.

Mail Limits: per day (Default: unlimited), per month (Default: unlimited)

Recipients types managed by LISTSERV	Enabled	Disabled	Hidden
Send to an Existing LISTSERV List (Recipients and Target Groups)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Let LISTSERV Select Recipients from a Database or LDAP Directory (Recipients and Target Groups)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

Note: No mail limit is applicable for the recipients types managed by LISTSERV.

Other	Enabled	Disabled	Hidden
Send to a Recipient Target Group (For the recipients types enabled above)	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

Note: The mail limit for target groups depends on the underlying recipients type.

In the lower section of the screen, set the recipient upload restrictions by entering a file name and path if the recipients will be taken from a file on the server "just-in-time" before the job is sent. If the list is left blank, no file access will be allowed. See the online help for more information on using this setting.

Figure 7-14 Recipients Restrictions - Upload Restrictions

Recipients Upload Restrictions

Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

Security Issue: Recipient upload of type "file" will access the files entered by the user in the context of the **server**. Therefore, in order to protect sensitive or other non-public information, you need to designate specific files and folders that users will be able to access. See the help page for more information.

Prefix-Strings of files to which access is allowed:

Default: No file access allowed

- **Tracking Restrictions** – This option defines the type of tracking that will be available, by default.
 - **Tracking Type Restrictions** – To enable a tracking type, select **Enable**. To disable, but still display the tracking type, select **Disabled**. To disable and hide a tracking type, select **Hidden**. To use the default setting, select **Use Default**.
- **Action Tracking Settings** – To enable action tracking, select **Enable Action Tracking**. To disable, select **Disable Action Tracking**.



Note: For information on setting the default tracking restrictions, see Section 5.2 [Application Default Settings](#).

Figure 7-15 Tracking Restrictions

Tracking Restrictions

Tracking Type Restrictions

Define which types of tracking are available by default. Select "Enabled" to enable a tracking type. Select "Disabled" to disable but still display a certain tracking type. Select "Hidden" to disable and hide a certain tracking type.

	Enabled	Disabled	Hidden
Personal Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Anonymous Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Unique Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Blind Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

Action Tracking Settings

Define the settings for cookie-based action tracking.

Enable Action Tracking
 Disable Action Tracking

OK Cancel



- **Hosted Data Settings** – Defines the settings for hosted recipient data.
 - **Hosted List Settings** – Select the hosted list type the group members can create.
 - If **Hosted LISTSERV Lists** will be used, optional suffixes for the LISTSERV lists can be used by selecting the option. Since there may not be more than one list with the same name on the same LISTSERV server, employing automatically generated or individual group/name suffixes for LISTSERV list names prevents list creation errors due to duplicate list names. See the online help for more details about suffixes.



Tip: By installing separate LISTSERV instances, and configuring the LISTSERV connection for each user/group to use a different LISTSERV instance, then the use of suffixes is not necessary to distinguish between lists. Separate instances will also give each group a separate LISTSERV Web Interface for their lists.



Note: The **Suffix** field is only visible if the **Use individual user/group name suffixes for LISTSERV Lists** option was selected while setting up the default hosted data settings. For more information, see Section 5.2 [Application Default Settings](#).



- **Hosted Data Import Restrictions** – Define the file import restrictions and prefix strings for the files that access is allowed for. Access to the server files is restricted and only files that start with one of the prefixes defined are allowed.



Important: In the User Interface, the option to import the data from a text file is only available in the Member/Subscriber Importer wizard if there is at least one such prefix defined. Similarly, the option to define an importer that reads from an external database is only available if the **Users/groups are allowed to import hosted data from an external database** option is checked.

A Recipient Importer can either pull data from a database or read it from a file. In the latter case, the file can either be provided by the user at the moment the Recipient Importer is launched, or it can be read from a pre-defined location on the server (defined during importer definition). In addition, a Recipient Importer can be launched manually (from inside User Interface) or externally (with an external trigger, via a URL). Depending on the importer type and the method that is used for launching it, the following will happen:

- The Recipient Importer reads from a database and is then launched manually. As soon as the importer is launched, it reads the current data from the external database and imports it into the hosted object.
- The Recipient Importer reads from a database and is then launched with an external trigger. As soon as the importer is launched, it reads the current data from the external database and imports it into the hosted object. (The request to the trigger URL does not contain any additional data, as the data is read from the database.)
- The Recipient Importer reads from a file that is provided by the user during launch, and then the importer is launched manually. When the importer is launched, the system first asks the user to select a file to upload. This is the file that the importer is supposed to process. Once the file is uploaded (from the user's local system), the importer processes it.
- The Recipient Importer reads from a file that is provided by the user during launch, and then the importer is launched with an external trigger. The request to the external trigger URL must be a special POST-request that already contains the file data that the importer will process. (That is, the POST-request to the trigger URL is two things at once – it is the actual trigger that starts the importer, and it is also the upload that provides the data for the importer.)
- The Recipient Importer reads from a file at a pre-defined server location, and then the importer is launched manually. When the importer is launched, it reads the file from the pre-defined location at the server and processes it. The location was pre-defined when the importer was initially created (in the Member/Subscriber Importer wizard).

- The Recipient Importer reads from a file at a pre-defined server location, and then the importer is launched with an external trigger. When the importer is launched, it reads the file from the pre-defined location at the server and processes it. The location was pre-defined when the importer was initially created (in the Member/Subscriber Importer wizard). (The request to the trigger URL does not contain any additional data, as the data is read from the file on the server.)
- **DomainKeys Settings** – Define whether or not DomainKeys signatures are used by default, and whether or not the users are allowed to change the default behavior for specific jobs. If defined on the application level, then these settings will be used as defaults for all accounts that do not have individual settings on a group or user level.

Figure 7-16 DomainKeys Settings

DomainKeys Settings

Define if e-mails that are sent by the user shall use DomainKeys signing by default. Additionally, define if the user is allowed to decide about DomainKeys signing (using the default value defined above).

Settings for Mail Jobs

Default Setting

Use inherited value: Yes, use DomainKeys signing

No, do not use DomainKeys signing

Yes, use DomainKeys signing

Mail Job Specific Settings

Use inherited value: The user may change the setting supplied above

The user must use the setting supplied above without changes for each mail job

The user may change the setting supplied above on a per-job basis

Settings for Hosted LISTSERV Lists

Use inherited value: Do not enforce DomainKeys-signing

Enforce DomainKeys-signing for all LISTSERV Maestro standard postings to Hosted LISTSERV Lists.

Do not enforce this and allow the user to disable or enable DomainKeys-signing on a per-list basis.

OK Cancel

Digitally signing email messages following the DomainKeys standard is a means to assert that the message originated from the domain that is claimed in the "From:" address. The digital signature is created for the whole message, which has the additional benefit that the recipient (once he or the receiving MTA has verified the signature) can be sure that the message has not been modified on its path from the sender to the recipient. Before enabling DomainKeys support in the application, bear in mind that if DomainKeys signatures are enabled for a mail job, then all messages from the mail job must be run through a signature computation, which in most cases slows down mail job delivery.

- **Settings for Mail Jobs:**
 - **Default Setting** – Defines the default behavior for DomainKeys signing.
 - **Mail Job Specific Settings** – Defines whether or not the default behavior for DomainKeys signing can be overridden for specific mail jobs.
- **Settings for Hosted LISTSERV Lists** – Define whether or not the default behavior for DomainKeys signing can be overridden for Hosted LISTSERV Lists.



Note: Changing the settings on this screen only applies to mail jobs that have not yet been authorized for delivery.

- **Auto-Archive Settings** – Defines whether or not jobs will be automatically archived when the completed job reaches the “auto-archive age”. If the **Auto-Archive Age** field is set to 0, then the auto-archive setting is off. Leave empty to use the default.

Figure 7-17 Auto-Archive Settings



- **External Access** – Defines whether or not the user account or group can export job data through an external access URL. If this Job Data Export feature is **Enabled**, then the user account or group can export a specific job or range of jobs, depending on the access parameters. The export is triggered by making a request to an external access URL with corresponding parameters and a security token. The result is an export file in XML format.

Figure 7-18 External Access Settings



Important: To protect against unauthorized access during the export process, the export request must include a secret security token. Every group, and every non-group account, has its own security token. The token is displayed on the Group Overview or User Account Overview screen once the Job Data Export feature is enabled.



Tip: For complete details on exporting your job data into an XML file, please see the online help for the External Access Setting screen, the [LISTSERV Maestro 4.0 What's New Manual](#), or the [Exporting LISERV Maestro Job Data to an XML File](#) whitepaper.

7.5 Deleting a User Account, Identity, or Group

To delete a user account, click on the account from the Accounts and Identities screen. The User Account Overview screen opens. From here, click on the **User Account** menu and select **Delete**. The Delete User Account screen opens. This screen allows you to delete the currently selected user account. Depending on what kind of a user account is currently selected one of the following options pairs will be available:

Account not in a group:

- **Keep data owned by the account:** Only the account itself and its personal settings will be deleted. The other data that is owned by the account will not be deleted but will remain in the system as "ownerless". The administrator can then later assign ownership of this data to a different user account or group or can delete the data separately.
- **Delete the account with all owned data and personal settings:** Together with the user account, all data and settings that are owned by or associated with this account will also be deleted.

Group-account, with other accounts existing in the same group:

- **Keep data owned by this account:** Only the account itself and its personal settings will be deleted. The other data that is owned by the account will not be deleted but will remain in the system as "ownerless". The administrator can then later assign ownership of this data to a different user account or group or can delete the data separately. All group owned data will not be deleted either.
- **Delete the account with all owned data and personal settings:** Together with the user account, all data that is owned by or associated with this account will also be deleted. Any group owned data will not be deleted.

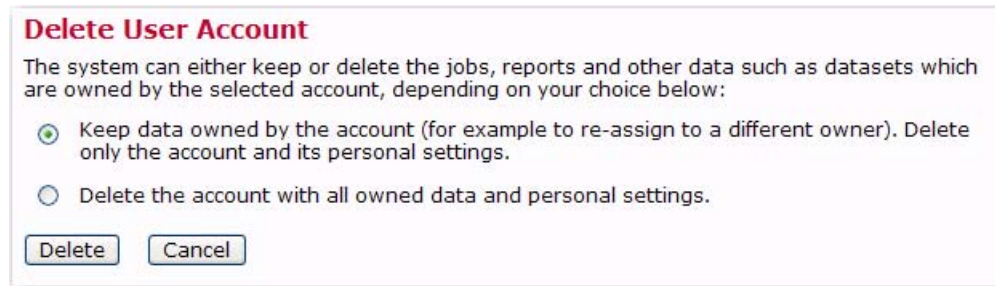
Group-account, last account in the group:

- **Keep data owned by the account or group:** Only the account itself and its personal settings will be deleted. The other data that is owned by the account or its group will not be deleted but will remain in the system as "ownerless". The administrator can then later assign ownership of this data to a different user account or group or can delete the data separately. The group specific settings will also not be deleted.
- **Delete the account with all owned data and personal settings:** Together with the user account, all data that is owned by the account or its group will be deleted, as well as the personal settings of the account and the group specific settings.

Click **[Delete]** to delete the currently selected account with the selected option, or click **[Cancel]** to leave the page without deleting the account.

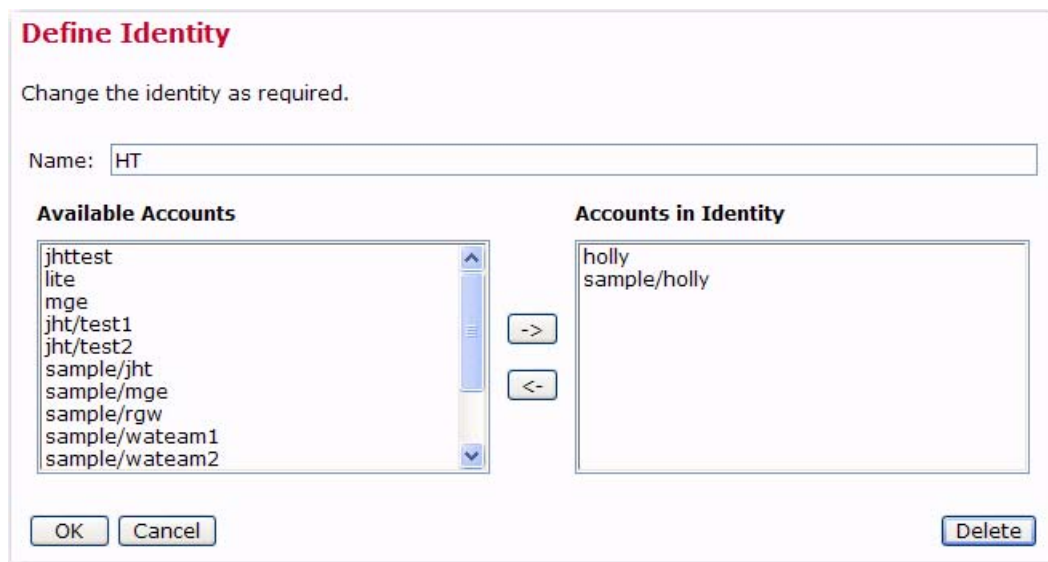
Any jobs, reports, sender profiles, and drop-in content elements owned by that user become “orphaned” and need to be reassigned. See Section 8 [Special Administrative User Account](#).

Figure 7-19 Deleting an Account



To delete an identity, select the identity from the Accounts and Identities screen. The Define Identity screen opens. From here, click on the **[Delete]** button.

Figure 7-20 Deleting an Identity



To delete an entire group, select the group from the Accounts and Identities screen. The Group Overview screen opens. Click on the **Group** menu and select **Delete**. The Delete Group screen opens.

This screen allows you to delete the currently selected user account using one of the following options:

- **Keep data owned by the group and its accounts:** Only the group specific settings and the accounts in the group with their personal settings will be deleted. The other data that is owned by the group or its accounts will not be deleted but will remain in the system as "ownerless".

This option gives the administrator the chance to later assign ownership of this data to a different user account or group or can delete the data separately.

- **Delete the group and all group accounts, with all owned data and settings:** Together with the group and all user accounts in it, all data and settings that are owned by or associated with this group or its accounts will also be deleted.

Click **[Delete]** to delete the currently selected group with the selected option, or click **[Cancel]** to leave the page without deleting the group.

Figure 7-21 Deleting a Group



Delete Group

The system can either keep or delete jobs, reports and other data such as datasets, depending on your choice below:

- Keep data owned by the group and its accounts (for example to re-assign to a different owner). Delete only the group settings and all group accounts with their personal settings.
- Delete the group and all group accounts, with all owned data and settings.

Delete Cancel

Section 8 Special Administrative User Account

With every installation of LISTSERV Maestro, a special user account for the system administrator is available. From this account, it is possible to archive delivered jobs and import jobs from the archive back into LISTSERV Maestro. The administrator can also change the ownership of a job, report, sender profile, drop-in content element, or recipients target group.



To access this account, simply log into the LISTSERV Maestro Administration HUB and use the **Switch to LUI** link at the top of the page. (For details regarding this logon feature, see Section 1.4 [Accessing the Hub Administrator Interface](#).)

8.1 The Toolbar

The Toolbar contains menus and icons that give you quick access to the different functions in LISTSERV Maestro. The menus may vary slightly for the administrator.




Figure 8-1 The Toolbar in the LISTSERV Maestro User Interface



The **Mail Job** menu lets you create a new job, view a list of open jobs (jobs that you have not finished setting up the criteria for), view a list of ongoing jobs (jobs that repeat), and a list of completed jobs. The options available vary depending on where you are in the application and what function you are performing. The possible options are:

- **New** – Create a new job.
- **Import** – Import a previously archived job from the archives.
- **Open Jobs** – Display the list of open jobs that have not been authorized for delivery.
- **Ongoing Jobs** – Display the list of ongoing jobs that are authorized for delivery and that are currently being processed or awaiting delivery.
- **Completed Jobs** – Display the list of completed jobs that have been delivered or closed after a delivery failure.
- **Archived Jobs** – Display the list of archived mail jobs.
- **Job Info** – Define the general information of the email job, such as the job title.
- **Job Journal** – View the job's history.
- **Copy Settings From Other Variant** – Copy the job definition settings from another variant job of the same A/B-split job to the current variant job. The copied settings will include the message definition, tracking definition, sender definition, and test delivery definition.
- **Authorize Delivery** – Authorize the job for delivery.
- **Authorize Sampling Variant Delivery** – Authorize the sampling variant jobs of an A/B-split job with Sampling for delivery.



- **Authorize Main Variant Delivery** – Authorize the main variant job of an A/B-split job with Sampling for delivery.
- **Team Collaboration** – Define the team collaboration settings for the job.
-  • **Trigger Job Delivery** – Trigger the job for delivery. If triggered, delivery will start immediately
-  • **Trigger Main Variant Delivery** – Trigger the main variant job of an A/B-split job with Sampling. If triggered, delivery will start immediately.
-  • **Trigger Delivery of Authorized A/B-Split Job Variants** – Trigger the A/B-split job variants for delivery. If triggered, delivery will start immediately
- **Revoke Send Authorization** – Revoke the delivery authorization of the job.
- **Revoke Authorized Variants of A/B-Split Jobs** – Revoke the delivery authorization of all authorized variant jobs of an A/B-split job.
- **Stop Sending** – Stop delivery of the job.
- **Re-Open Job** – Re-open the job for further editing.
- **Re-Open Failed Variants of A/B-Split Job** – Re-open the failed variant jobs of an A/B-Split job for further editing.
- **Close Job** – Close the job in its failed state.
- **Close All Open Variants of A/B-Split Job As Failed** – Close the open variant jobs of an A/B-split job as failed.
- **Retry Sending** – Retry the delivery of the failed job.
- **Resume Sending** – Resume the delivery of the job to the pending recipients.
- **Comparison Report for Completed Variant Jobs** – Generate the Job Comparison report for the selected A/B-split testing variant jobs.
- **Move Job to Archive** – Move the selected mail job to the archive. An archived job will no longer be accessible in the normal LISTSERV Maestro User Interface; it will appear as if the job was deleted. Only the LISTSERV Maestro administrator can import an archived job. Also, an imported job will be “frozen” into the state that it had when it was initially archived (for example, no more tracking events will be collected for this job).
- **Delete Job** – Delete the job.
- **Multiple Job Actions** – Work with one or several jobs at once with one of the following options:
 - **Job Comparison Report** – Generate the Job Comparison report for the selected jobs.
 - **Edit Category of All Selected Jobs** – Edit the job category of all selected jobs.
 - **Edit Team Collaboration of All Selected Jobs** – Edit the team collaboration settings of all selected jobs.
 - **Move All Selected Jobs to Archive** – Move all selected jobs to the archive.
 - **Delete All Selected Jobs** – Delete all selected jobs (not recoverable).

- **Select All** – Select all jobs in the current job list.
- **Unselect All** – Unselect all jobs in the current job list.
- **Enable Multiple Job Actions** – Enable the multiple job actions. Checkboxes will appear next to every job, allowing multiple job selection, and the **Multiple Job Actions** menu will be displayed.
- **Disable Multiple Job Actions** – Disable the multiple job actions. Checkboxes and the **Multiple Job Actions** menu will be hidden.
- **Change Owner of All Selected Jobs** – Change the owner of all selected jobs.
- **Import all Selected Jobs** – Import all selected jobs from the archive.



Notes: By default, multiple job actions are disabled and the menu contains only the **Enable Multiple Job Actions** option. By selecting this option, multiple job actions become enabled, which has two effects: The menu now contains more options (with various actions that can be applied to several jobs at once), and in the job list, a checkbox is shown in front of each listed job. Check those jobs that you want to manipulate, then select the desired option from the **Multiple Job Actions** menu. This action will then be applied to all selected jobs. For the administrator, the multiple job actions are always enabled and can not be disabled.

The **Report** menu lets you create a new report, view any existing reports, and open the Delivered Jobs Statistics and Delivered Recipients Statistics reports. The options available vary depending on where you are in the application and what function you are performing. The possible options are:

- **New Report** – Create a new tracking statistic report.
- **Reports** – Display the list of currently available tracking statistics reports.
- **Recipient Action Tracking** – Activate/deactivate recipient action tracking and define action tags.
- **Completed Jobs Statistics** – Display the statistical report about the previously delivered jobs.
- **Delivered Recipients Statistics** – Display the statistical report about the recipients of previous jobs.
- **System Metrics Reports** – For the administrator, display the System Metric Reports.
- **Add Report to Dashboard** – Add the displayed report to the Dashboard.
- **Apply Settings to Report on Dashboard** – Apply report changes to the report shown on the Dashboard.

The **Recipient Warehouse** menu lets you create new and manage datasets, lookup tables, and target groups. You can also edit the bounce process settings for the warehouse. The options available vary depending on where you are in the application and what function you are performing. The possible options are:

- **New Dataset** – Create a new recipient dataset.
- **Datasets** – Display a list of currently available recipient datasets.

- **Datasets & Lookup Tables** – Display the administrator’s page for recipient datasets and lookup tables.
- **New Lookup Table** – Create a new lookup table.
- **Lookup Tables** – Display the list of currently available lookup tables.
- **New Recipient Target Group** – Create a new target group using one of the following options:



- **Based on Dataset** – Create a new target group that is based on a recipient dataset.
- **Based on Hosted Recipient List** – Create a new target group that is based on a hosted recipient list.
- **Based on Hosted or Standard LISTSERV List** – Create a new target group that is based on a classic LISTSERV list.
- **Based on Uploaded Text File** – Create a new target group that is based on an uploaded text file (CSV-file).
- **Based on Database Access by LISTSERV Maestro** – Create a new target group based on a database accessed by LISTSERV Maestro.
- **Based on Database Access by LISTSERV** – Create a new target group based on a database accessed by LISTSERV.
- **Target Groups** – Display the list of currently available target groups.
- **Bounce Processing** – Define the bounce processing strategy for the Recipient Warehouse.

The **Utility** menu lets you create new and manage drop-in content and sender profiles; plus, it lets you change your account password. The options available vary depending on where you are in the application and what function you are performing. The possible options are:

- **New Content Template** – Sub-menu for creating new content templates:
 - **Create Empty Template** – Create a new (empty) content template.
 - **Create Template Based on Mail Job** – Create a new content template that is based on the message definition of an existing mail job.
 - **Upload Template** – Create a new content template by uploading a content template definition file (ZIP-file).
- **Content Templates** – Show the list of currently available content templates.
- **New Drop-In Content** – Create a new drop-in.
- **Drop-Ins** – Display the list of currently available drop-ins.
- **New Sender Profile** – Create a new sender profile.
- **Sender Profiles** – Display the list of currently available sender profiles.
- **Change Password** – Change the password of the current user account.
- **Print Page** – Print the current page.



- **About LISTSERV Maestro** – Show detailed information about the current LISTSERV Maestro version.

The **Dashboard** menu lets you create new and manage the configuration of your Dashboard(s). This menu is only available from the Dashboard screen. The available options are:

- **Switch Dashboard** – Show the list of available Dashboards. From this list, select which one will be displayed as the current Dashboard in the current session. Once selected, that particular Dashboard will be displayed.
- **Hide/Show Default Sections** – Hide or show any of the default Dashboard sections (applies only to the currently selected Dashboard).
- **Edit Dashboard Name** – Edit the name (title) of the currently selected Dashboard.
- **Create New Dashboard** – Create a new Dashboard with all default sections, in default order, without any Dashboard reports.
- **Copy Current Dashboard** – Create a new Dashboard as a copy of the current Dashboard.
- **Delete Current Dashboard** – Delete the current Dashboard.

The **Dataset** menu lets you manage the datasets, categories, and members. The options available vary depending on where you are in the application and what function you are performing. The possible options are:

- **Dataset Overview** – Return to the Recipient Dataset Details screen.
- **Edit Dataset Settings** – Edit the dataset settings.
- **Edit Category** – Edit the settings of the category that is currently selected in the dataset tree.
- **Create** – Create hosted lists and categories using one of the following options:
 - **Create Hosted Recipient List** – Create a new Hosted Recipient List in this dataset.
 - **Create Hosted LISTSERV List** – Create a new Hosted LISTSERV List in this dataset.
 - **Create Category** – Create a new category (either in the currently selected category or in the root of the dataset tree).
 - **Convert or Clone Standard LISTSERV List** – Convert an existing, non-hosted LISTSERV list into a Hosted LISTSERV List or creates a Hosted LISTSERV List as a clone of an existing non-hosted LISTSERV list without changing the original list.
- **Browse/Edit Confirmed Members** – Browse through the members of the dataset and edit them manually.
- **Browse Confirmed Members** – Browse through the members of the dataset.
- **Add Single Member** – Add a single, new member to the dataset.



- **Member Import** – Add, update, or delete members in the dataset by uploading or importing a file.

- **Import Members Now** – Open the Import Members wizard where you can define which members to import at this moment.
- **New Member Importer** – Define member importer for future use.
- **Member Importers** – Display list of member importers for the selected dataset.
- **Download all Members** – Download the current members of the dataset.
- **Team Collaboration** – Define the team collaboration settings for the dataset.
- **Refresh and Manage Linked LISTSERV Lists** – Define which of the non-hosted LISTSERV lists at the LISTSERV server to include in the dataset as Linked LISTSERV Lists. This option will also refresh the settings of the already linked LISTSERV lists by reading their current settings (as defined in the list headers) from LISTSERV so that LISTSERV Maestro will be aware of any changes to the lists.
- **Subscriber Access URLs** – Display the URLs for the Subscriber Access pages of the dataset's member area.
- **Tree Structure** – Manage operations in a dataset tree using one of the following options:
 - **Cut** – Cut the selected category from the dataset tree in preparation for a Cut & Paste operation (more precisely, the node will be marked as “cut”, but it will only be removed from its current parent node once you select **Paste** in another node).
 - **Copy (Categories Only)** – Copy the selected category (and its subcategories) in the dataset tree in preparation for a Copy & Paste operation (more precisely, the category will be marked as “copied”, but it will only be created once you select **Paste** in another node).



Note: Any hosted lists in the category or sub category will not be copied.

- **Paste** – Paste a node into the currently selected node or root of the dataset tree. The node to paste must have been marked for “cut” or “copy” (see the **Cut** and **Copy** options above).
- **Clear Cut/Copy State** – Clear the current “cut” or “copy” state and unmarks the node that was to be “cut” or “copied” (see the **Cut** and **Copy** options above).
- **Delete all Members** – Delete all members (including unconfirmed ones) from the dataset.
- **Delete Dataset** – Delete the dataset, including all members and lists contained within it.
- **Delete Category** – Delete the category that is currently selected in the dataset tree.

The **Hosted List** menu contains options that are related to the hosted list currently selected in the dataset tree. The options available vary depending on where you are in the application and what function you are performing. The possible options are:

- **Edit List Details** – Edit the hosted list.
- **Browse/Edit Confirmed Subscribers** – Browse through the subscribers of the hosted list and edit them individually.



- **Browse Confirmed Subscribers** – Browse through the subscribers of the hosted list (view only).
- **Add Single Subscriber** – Add a single, new subscriber to the hosted list.
- **Subscriber Import** – Add, update, and delete subscribers in the hosted list by uploading or importing a file.
 - **Import Subscribers Now** – Open the Import Subscribers wizard where you can define which subscribers to import at this moment.
 - **New Subscriber Importer** – Define subscriber importer for future use.
 - **Subscriber Importers** – Display list of subscriber importers for the selected dataset.
- **Download all Subscribers** – Download the current subscribers of the hosted list.
- **Delete all Subscribers** – Delete all subscribers from the hosted list.
- **Delete List** – Delete the hosted list with all subscribers.

The **Linked LISTSERV List** menu contains options that are related to the Linked LISTSERV List currently selected in the dataset tree. The options available vary depending on where you are in the application and what function you are performing. The possible options are:

- **List Configuration** – Open the corresponding LISTSERV Web Interface page for management of the selected Linked LISTSERV List.
- **Convert To or Clone As Hosted List** – Convert the selected Linked LISTSERV List into a Hosted LISTSERV List or clone a Hosted LISTSERV List from it (cloning will not change the original Linked LISTSERV List).
- **Refresh and Manage Linked LISTSERV Lists** – Define which of the non-hosted LISTSERV lists at the LISTSERV server to include in the dataset as Linked LISTSERV Lists.
- **Remove Link to LISTSERV List** – Remove the link to this LISTSERV List (removes the Linked LISTSERV List node from the dataset tree and the actual list at LISTSERV will remain the unchanged).



The **Lookup Table** menu contain options for the currently selected lookup table. The possible options are:

- **Lookup Table Settings** – Edit general lookup table settings (name, description, encoding, secondary columns).
- **View References** – View the dataset and hosted list profile fields that reference the lookup table.
- **Add New Entry** – Add a new entry to the lookup table.
- **Upload Entries** – Upload new entries into the lookup table.
- **Download Entries** – Download the lookup table entries.
- **Define New Subset** – Define a subset of lookup table entries.
- **Edit Highlighted Subset** – Edit the currently highlighted lookup table subset.

- **Undefine Highlighted Subset** – Undefine the currently highlighted lookup table subset.
- **Delete All Entries** – Delete all entries from the lookup table (but not the lookup table itself).
- **Delete Lookup Table** – Delete the lookup table.

The **LISTSERV** menu takes you to specific areas in the LISTSERV Web Interface. The options available vary depending on where you are in the application and what function you are performing. The possible options are:

- **Server Administration Dashboard** – Open the Server Administration Dashboard of the LISTSERV Web Interface.
- **List Administration Dashboard** – Open the List Administration Dashboard of the LISTSERV Web Interface.
- **Site Configuration** – Open the site configuration page of the LISTSERV Web Interface.
- **LISTSERV Archives** – Open the LISTSERV archives page of the LISTSERV Web Interface.
- **Server Reports** – Contains options for reports in the LISTSERV Web Interface.
 - **Server Usage Reports** – Open the server usage reports page of the LISTSERV Web Interface.
 - **Anti-Virus & Anti-Spam Reports** – Open the anti-virus and anti-spam reports page of the LISTSERV Web Interface.
 - **Server Activity Reports** – Open the server activity reports page of the LISTSERV Web Interface.

The **Back To** menu is a quick way to return to the previous screen or to one of the recently visited screens. The possible options are:

- **Dashboard** – Go back to the Dashboard screen.
- **Selected Job "JOB NAME HERE"** – Go back to the job with the given name.
- **Selected A/B-Split Job "A/B-SPLIT JOB NAME HERE"** – Go back to the A/B-split job with the given name.
- **Create New Job** – Go back to the Start New Job screen.
- **Selected Report "REPORT NAME HERE"** – Go back to the tracking report with the given name.
- **Create New Report** – Go back to the Define Report screen.
- **Job Comparison Report** – Go back to the Job Comparison Report screen.
- **Selected Dataset "DATASET NAME HERE"** – Go back to the dataset with the given name.
- **Create New Dataset** – Go back to the Recipient Dataset Definition wizard.
- **Selected Lookup Table "LOOKUP TABLE NAME HERE"** – Go back to the lookup table with the given name.
- **Create New Lookup Table** – Go back to the Lookup Table screen.

- **Selected Target Group "TARGET GROUP NAME HERE"** – Go back to the target group with the given name.
- **Create New Target Group** – Go back to the Target Group Definition wizard.
- **Selected Content Template "TEMPLATE NAME HERE"** – Go back to the content template with the given name.
- **Create New Content Template** – Go back to the Edit Content Template screen.
- **Selected Drop-In "DROP-IN NAME HERE"** – Go back to the drop-in with the given name.
- **Create New Drop-In** – Go back to the Drop-In Content Element screen.
- **Selected Sender Profile "SENDER PROFILE NAME HERE"** – Go back to the sender profile with the given name.
- **Create New Sender Profile** – Go back to the Sender Profile screen.

The **Preferences** menu contains options for defining various user preferences. The possible options are:

- **User Interface** – Contains options for defining preferences for the User Interface.
 - **Date and Number Formats** – Define the date, time, and number preferences for the User Interface.
 - **Display** – Define the display preferences for the User Interface.
- **Mail Jobs** – Contains options for defining job-related preferences.
 - **New Mail Job** – Define the preferences for a new job, including the team collaboration preferences.
 - **Recipients** – Define the preferences for the recipient definition of a job.
 - **Content** – Define the preferences for the content definition of a job.
 - **Tracking** – Define the preferences for the tracking definition of a job.
 - **Delivery Test** – Define the preferences for the delivery test of a job.
 - **Scheduling** – Define the preferences for the schedule definition of a job.
- **Reports** – Contains options for defining preferences for reports.
 - **New Report** – Define the general tracking report preferences.
- **Datasets** – Contains options for defining preferences for datasets.
 - **Team Collaboration** – Define preferences for the team collaboration settings of a recipient dataset.

The **Logout** menu logs you out of the LISTSERV Maestro Interface. For a user account that is not part of an identity, then this menu appears as a single item, which, if you click on it, allows you to log out from LISTSERV Maestro. For a user account that is part of an identity, then this menu has the following options:

- **Logout this Account** – Log out from LISTSERV Maestro.
- **Switch Account** – Switch to a different account in the same identity (without the need to logout and login again).



The **Help** icon is used to access the help associated with the current screen.



8.1.1 Behavioral Changes to the Toolbar

The behavior of the **Dataset** menu on the Toolbar has been updated to accommodate those who may need to access this menu while working with lists. Because of this, there are several changes to the Toolbar:

- As soon as you select a dataset, the **Dataset** menu is displayed. This menu remains visible the entire time the dataset is selected. Therefore, if you select one of the Hosted Recipient List (HRL), Hosted LISTSERV List (HLL), or Linked LISTSERV List (LLL) nodes in the dataset tree, then the **Dataset** menu is still displayed.
- If you select a HRL/HLL node in the dataset tree, then the **Hosted List** menu is displayed in addition to the **Dataset** menu.
- If you select a LLL node in the dataset tree, then the **Linked LISTSERV List** menu is displayed in addition to the **Dataset** menu.
- If you select a category node (folder) in the dataset tree, then, as before, only the **Dataset** menu is visible. However, previously, this **Dataset** menu only showed the category/folder related menu items. This behavior has been changed. If a category is selected, then the **Dataset** menu now continues to show the full set of dataset related menu items, and the category specific items are shown in addition to the other menu items.
- In the **Hosted List** menu, the following menu items have been removed, as they were duplicates of menu items that already exist in the **Dataset** menu (they are no longer required in the **Hosted List** menu because the **Dataset** menu is now always visible):
 - **Dataset Overview**
 - **Tree Structure** (the whole sub-menu has been removed)
- Similarly, in the **Linked LISTSERV List** menu, the following duplicate menu items have been removed:
 - **Dataset Overview**
 - **Refresh and Manage Linked LISTSERV Lists**
 - **Tree Structure** (the whole sub-menu has been removed)

8.2 The Dashboard

The opening screen of LISTSERV Maestro's interface is called the Dashboard, providing a quick summary of what has happened recently and what is due to happen in the future.

The Dashboard contains five sections that may be displayed, hidden, or rearranged. This allows for a convenient and easy way to customize the information shown, making it possible for you to see what is most important to you. The Dashboard sections are:



- **System Metrics** – This section gives you a complete overview of every group and account for a specific period of time.
- **Currently in the System** – This section contains the general statistics for your account, such as jobs in progress, number of jobs completed, etc.

- **Jobs Due Next** – This section contains a list of jobs that are currently open and that are due next for authorization. This section does not show any jobs that have a “Authorize Due By” date configured.
- **Current and Upcoming Deliveries** – This section contains a list of jobs that are currently being delivered and that are scheduled for an upcoming delivery. From this section, you can also view any failed jobs that have yet to be closed or re-opened.
- **Recent Deliveries** – This section contains a list of jobs that have been recently delivered and jobs that have attempted delivery but failed.



Note: The administrator has the option to show a summary of all user accounts, certain user accounts, or just your user account on the Dashboard.

Figure 8-2 The Dashboard

Dashboard Data retrieved at Oct. 24, 2008 11:48:21. [Refresh](#)

Show dashboard for: <All Accounts and Groups>

System Metrics x

System metrics for period: **Oct. 23, 2008 12:00:00 to Oct. 24, 2008 12:00:00**

Disk usage at the end of the period				Request metrics during period					
	LUI	TRK	HUB	Database		LUI	LIST	TRK	HUB
Used	200.8 MB	34.3 MB	10.7 MB	?	Total Requests	0	0	607	0
Free	110,254.2 MB	110,254.2 MB	110,254.2 MB		Average Response Time	?	?	15.2 ms	?

Objects in the system at the end of the period				Other metrics			
Open Jobs:	56	Reports:	8	Datasets:	18	Jobs sent during period:	1
Ongoing Jobs:	9	Tracking Events:	59,588	Hosted Lists:	20	Recipients sent during period:	281
Completed Jobs:	124			Hosted Recipients:	12,977	Duration of last backup:	1.5 min

Currently In The System x

<u>Open Jobs:</u>	56	Directly Distributed Recipients:	47,305	<u>Datasets:</u>	18
<u>Ongoing Jobs:</u>	9	Postings To LISTSERV Lists:	1	Hosted Lists:	20
<u>Completed Jobs:</u>	124	<u>Reports:</u>	8	Hosted Recipients:	12,977
of which tracked jobs:	119	Tracking Events:	59,588		

Current And Upcoming Deliveries x

Waiting for Delivery Trigger **080617A: Kayaking Newsletter**

Failed **080716B: Send to Dataset**

Failed **080716C: Test for Forward to a Friend tracking events**

[more...](#) (show up to jobs)

Recent Deliveries x

23 hours ago - at Oct. 23, 2008 12:51:31 - **DEMO-081022A: Anon**

47 hours ago - at Oct. 22, 2008 12:51:31 - **DEMO-081021A: Anon**

3 days ago - at Oct. 21, 2008 12:51:31 - **DEMO-081020A: Anon**

[more...](#) (show up to jobs)

Jobs Due Next x

Overdue - at May 4, 2006 07:13:00 - **060504C: This job is overdue!**

Overdue - at Aug. 21, 2006 15:00:00 - **CLICK-060821D: An overdue job**

Overdue - at Sep. 25, 2006 10:00:00 - **CLICK-060821C: A Job with a Due Date**

[more...](#) (show up to jobs)

8.2.1 Customizing the Dashboard

In LISTSERV Maestro, you can have several independent Dashboard configurations. Each Dashboard can contain its own unique set of default sections and/or reports. These sections can be easily rearranged and added so that the most important information appears at the top. For example, perhaps you are interested in seeing an overview of all jobs and subscriptions; maybe you want to focus on a particular category of jobs; or perhaps you would like to monitor the activities related to a particular dataset or even a single list. With the LISTSERV Maestro Dashboard, all of this is possible. In fact, you can even add reports to the Dashboard, giving you immediate access to your most important report statistics.

8.2.1.1 Updating the Dashboard's Data

To update the data displayed on the Dashboard, click the **Refresh** link. The date and time of the data being viewed is shown next to this link.

Figure 8-3 Refreshing the Dashboard



Note: The Refresh link does not update the data for any Dashboard reports. Dashboard reports will automatically update their data in regular intervals.

8.2.1.2 Editing the Dashboard's Name

To edit the name of the Dashboard currently being displayed, click on the **Dashboard** menu, and then select **Edit Dashboard Name**. The Edit Dashboard Name screen opens. Enter the new name for the Dashboard and click **[OK]**.

Figure 8-4 The Edit the Dashboard Name Screen

The screenshot shows a dialog box titled 'Edit Dashboard Name'. It contains a text input field with the text 'Weekly Newsletter Dashboard' and two buttons: 'OK' and 'Cancel'.

Changing the name of the Dashboard is a great way to keep track of what each Dashboard represents. For information on creating multiple Dashboards, see Section 8.2.1 [Customizing the Dashboard](#).

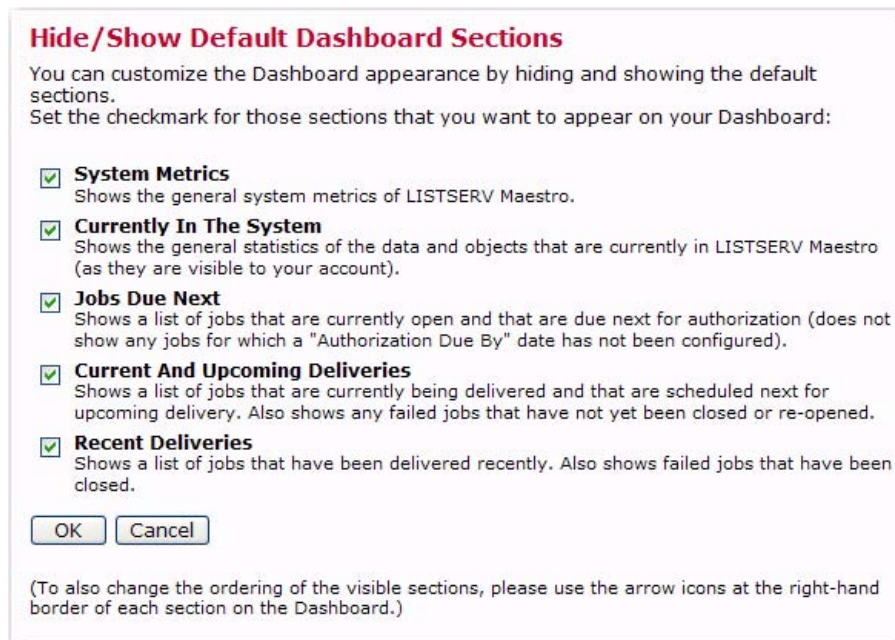
8.2.1.3 Hiding/Showing Dashboard Sections

When using the LISTSERV Maestro Dashboard, you'll notice that it is divided into multiple sections that can be re-organized quite easily. For instance, you can move any section up or down in order of importance, hide sections that don't interest you, and show only those that do. Changes to the Dashboard are constant from session to session, which means that any rearranging you do or reports that you add will be available each time you log in.

Each Dashboard section can be hidden or moved. The "x" box in the top right corner of each section will hide that particular section. The up and down arrows in the top and bottom right corners will cause the section to trade places with the section above or below it.

In addition, to display or hide a section, you can click on the **Dashboard** menu, and then select **Hide/Show Default Sections**. The Hide/Show Dashboard Sections screen opens.

Figure 8-5 The Hide/Show Default Section Screen



8.2.2 System Metrics

The administrator can now access and monitor system metrics in the User Interface.



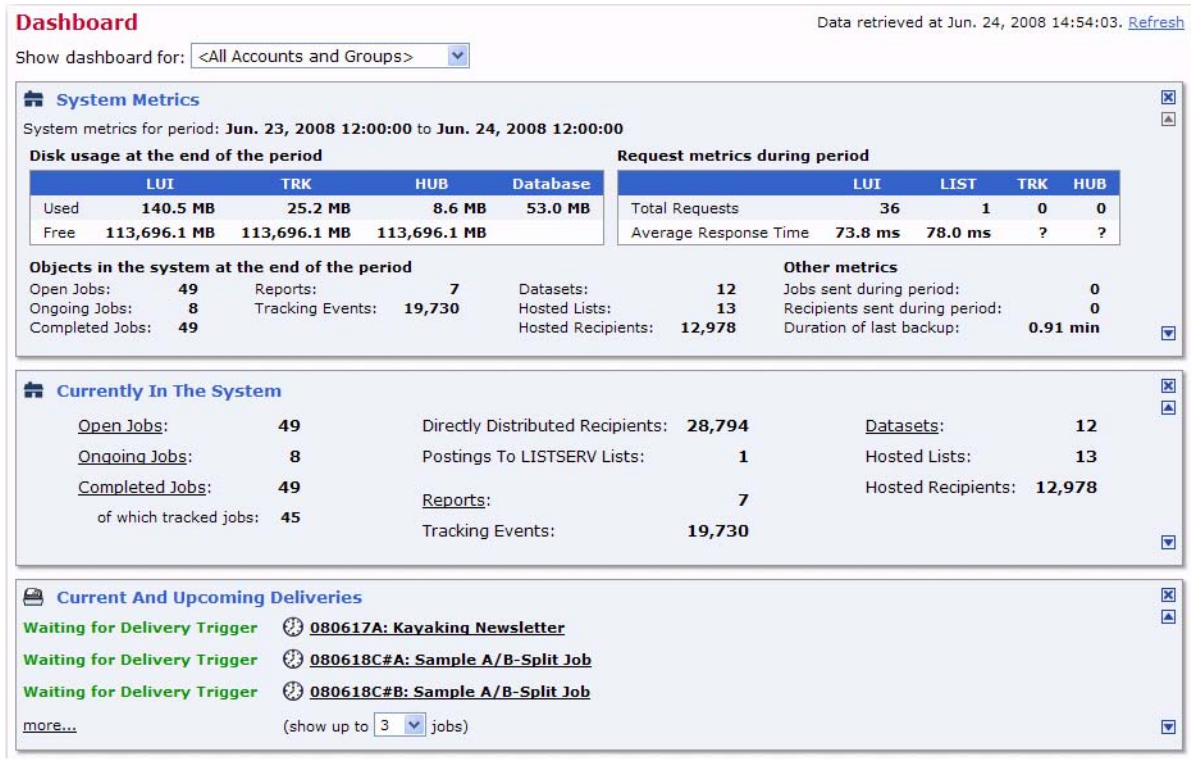
Note: Additional settings concerning this feature are available in the Administration HUB. For details, see Section 5.1 [Application Settings](#).

The System Metrics section gives you a complete overview of every group and account for a specific period of time. You can view:

- How much space the User Interface, the HUB, the Maestro Tracker, and your database are using, as well as how much space is free.

- How many objects are in the system, such as open, ongoing, and completed jobs; reports; tracking events; datasets; hosted lists; and recipients.
- How many jobs were sent and how recipients these jobs were sent to.
- How long the last backup took.

Figure 8-6 System Metrics on the Dashboard

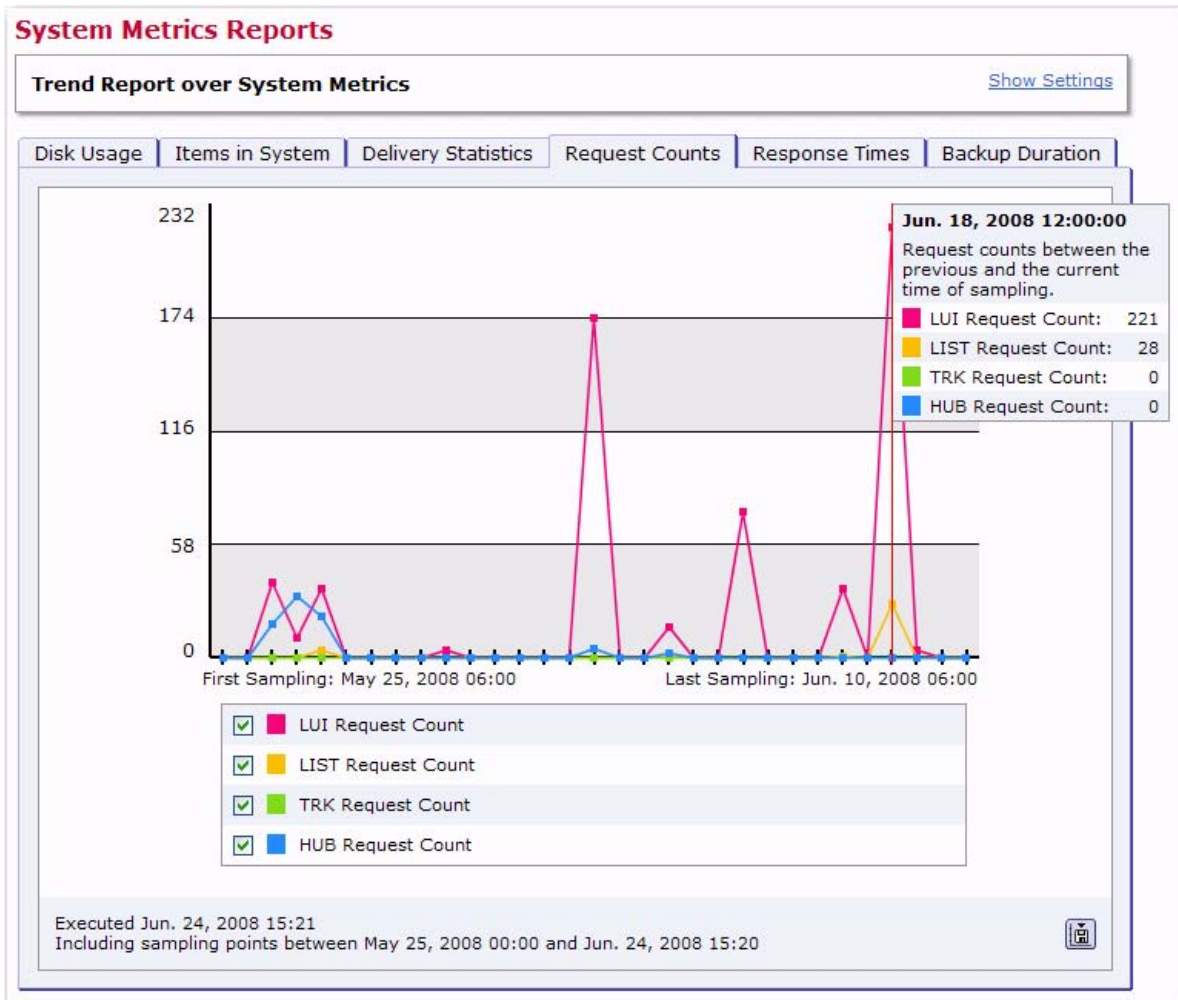


8.2.3 System Metrics Reports

Your system metrics can also be shown in a set of reports that show trends in how your system is running. Similar to other LISTSERV Maestro reports, you can drill-down into specific events on these reports to see detailed information, and you can also run them for the immediate past or for a fixed period of time.

To run a system metrics report, click on the **Reports** menu and select **System Metrics**. The System Metrics report opens with several tabs. Each tab contains a different report pertaining to a specific type of data.

Figure 8-7 The System Metrics Report - Request Counts Tab



8.2.3.1 Adding to the Dashboard

To add a specific System Metrics report to the Dashboard, click on the **Report** menu and select **Add Report to Dashboard**. This will only add the report that was being viewed (i.e. the report on the Request Counts tab); it does not add the entire set of reports. You can add up to 20 reports on the Dashboard; therefore, if you'd like to add other System Metric reports, simply click on their tab and then add it to the Dashboard.

If you click on the report on the Dashboard, the entire set of System Metrics reports will open. If you make any changes to the report that was on the Dashboard, then the **Apply Settings to Report on Dashboard** option will be available on the **Report** menu.

However, if you go to the System Metrics report by selecting it from the **Report** menu, and click on the tab to view that same report, then this option will not appear since this is a different version than the one on the Dashboard. If you'd like to add another version of that report to the Dashboard, then simply add it to the Dashboard using the **Add Report to Dashboard** option.



Note: The **Apply Settings to Report on Dashboard** option will only appear for that specific report that was on the Dashboard. For instance, if you added the Request Counts report to the Dashboard, then that is the only report/tab in the set of System Metrics reports that will have this option available.

8.2.3.2 Downloading

You can download the entire set of System Metrics reports or just a single report.



To download the entire set of reports, click on the **Download** icon from any tab. A pop-up menu appears with options to download all reports **as PDF** or **as Text File (CSV)**.

To download a specific report, go to that tab, and then click on the **Download** icon. A pop-up menu appears with an option to download that specific report **as PDF** or **as Text File (CSV)**.

Figure 8-8 Downloading the System Metrics Report



Tips on Download Formats:

- **Download as PDF** – The report will be downloaded in the form of a PDF file for easy distribution or printing.
- **Download as Text File (CSV)** – The report will be downloaded in the form of a ZIP-file that contains the current report in the form of a text file with comma separated values (a CSV-file), plus an additional readme.txt file with details about the report. This CSV-data can be imported into 3rd party tools (such as Excel) for further analysis.

8.2.4 Managing Dashboard Reports

Dashboard reports are for users who want immediate access to key email reports and statistics at a single glance, helping you make better decisions for any future email communication initiatives. Once a report is placed on the Dashboard, you can drill down

to problem areas, share success stories and best practices with your team, and quickly adjust your email campaigns. In addition, with one click, you have full access to the report for closer inspection.

LISTSERV Maestro gives you the ability to add up to 20 reports to the Dashboard. Each report is displayed in its own section, showing a "miniaturized" version of the report, and each report is automatically updated when new data is available.

8.2.4.1 Adding Reports to the Dashboard

To add a report to the Dashboard, go to and view the actual report that you want to add, and then click on the **Report** menu and select **Add Report To Dashboard**.

Click **[OK]** to add the report. If you want to view the added report on the Dashboard, click the **Back To** menu, and then select **Dashboard**. The report section will be located at the bottom of the Dashboard.

8.2.4.2 Re-Organizing Dashboard Reports

To move the report, you can use the up and down arrows. If you add another report to the Dashboard and move it up to the location directly below the first report, then the two reports will be displayed side by side. Left/right arrows are displayed for each report, giving you the ability to move the reports left and right. If you move the report up or down, each report will be on a line by itself and these new arrows will no longer be available.

To remove a report from the Dashboard, simply click the "x" box located in the upper right corner of that report section.

8.2.4.3 Viewing Additional Report Details


To view more details on the report, simply hover your mouse over the report.

To view the full report with all details on its own screen, click on the report. Once you have opened the Dashboard report, you can change the view settings; however, these changes are only temporary and only apply to the current report details. To apply these changes to the report on the Dashboard as well, click on the **Report** menu, and then select **Apply Settings to Report on Dashboard**.

8.2.4.4 Renaming the Dashboard Report

To rename the report on the Dashboard, click on the report title, enter the new name in the text box, and then click **[OK]**.

Figure 8-9 Renaming the Dashboard Report



Click **[OK]** to confirm the deletion of the current Dashboard.

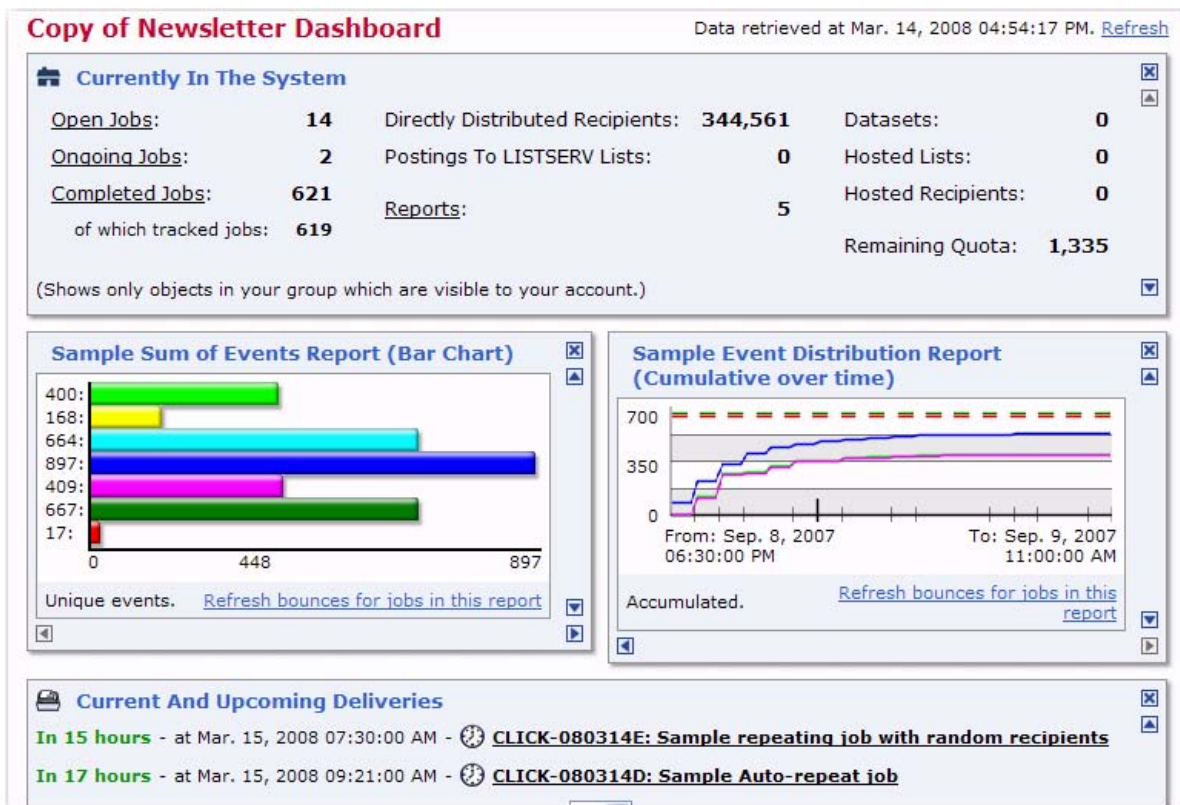
8.2.5 Creating Multiple Dashboards

Users can define and configure several Dashboards, each with a unique collection of standard sections and reports. Once these Dashboards are created, a user can easily switch back and forth between these Dashboards.

8.2.5.1 Copying the Current Dashboard

After customizing the Dashboard, you can create a new Dashboard based on your changes. To do so, click on the **Dashboard** menu and select **Copy Current Dashboard**. The copied Dashboard is displayed.

Figure 8-10 Copying a Dashboard



To rename the Dashboard, simply use the **Edit Dashboard Name** option on the **Dashboard** menu. For details, see Section 8.2.1.2 [Editing the Dashboard's Name](#).

8.2.5.2 Creating a New Dashboard from Scratch

To create a brand new Dashboard from scratch, click the **Dashboard** menu and select **Create New Dashboard**. The new Dashboard is displayed with the default sections.

Figure 8-11 Creating a Dashboard

The screenshot shows a dashboard titled "Dashboard (2)" with a "Refresh" link and a timestamp "Data retrieved at Mar. 13, 2008 08:46:19 AM". The dashboard is divided into five sections:

- Currently In The System:** A grid of metrics including Open Jobs (1), Ongoing Jobs (1), Completed Jobs (5), Directly Distributed Recipients (158), Postings To LISTSERV Lists (1), Reports (1), Tracking Events (1,157), Datasets (1), Hosted Lists (1), and Hosted Recipients (38).
- Recently Visited:** A section indicating "No objects visited so far." with a note "(Shows only objects visited in this session.)".
- Jobs Due Next:** A section stating "Currently there are no jobs due (or no jobs with a due date configured)."
- Current And Upcoming Deliveries:** A section showing "In 48 hours - at Mar. 15, 2008 08:00:00 AM - 070619A: Weekly Newsletter" with a "show up to 3 jobs" dropdown.
- Recent Deliveries:** A section listing three recent deliveries: "533 days ago - at Sep. 27, 2006 08:00:00 AM - 060827A: Test Job", "554 days ago - at Sep. 6, 2006 11:15:30 AM - 060906A: Newsletter", and "564 days ago - at Aug. 27, 2006 08:00:00 AM - 060810D: Test Job", with a "more..." link and a "show up to 3 jobs" dropdown.

For details on customizing the newly created Dashboard, see Section 8.2.1 [Customizing the Dashboard](#).

8.2.6 Switching Dashboards

If you have more than one Dashboard and would like to display a different Dashboard other than the one currently displayed, then you can simply click on the **Dashboard** menu and select **Switch Dashboards**. The Switch Dashboard screen opens.

Figure 8-12 Switching Dashboards

Switch Dashboard

Click the name of the dashboard that you want to use as the current dashboard, or click **[Cancel]** to leave the page without changing the current dashboard. (The current dashboard is emphasized in the list below.)

Dashboard	Standard Sections: Currently In The System, Recently Visited, Jobs Due Next, Recent Deliveries, Current And Upcoming Deliveries Dashboard Reports: Sample
Weekly Newsletter Dashboard	Standard Sections: Currently In The System, Recently Visited, Jobs Due Next, Current And Upcoming Deliveries, Recent Deliveries Dashboard Reports: <none>

From this screen, select the Dashboard you want to use in your current LISTSERV Maestro session. The currently selected Dashboard is emphasized in the list of available Dashboards. To select a different Dashboard, simply click on the name of the Dashboard.



Note: The choice you make here will only apply to the current session. This means that any parallel sessions that you may have open in other LISTSERV Maestro windows can have their own current Dashboard displayed.

To cancel and return to the current Dashboard, click **[Cancel]**.

8.2.7 Deleting Dashboards

If you have multiple Dashboards and would like to delete the currently displayed Dashboard, click on the **Dashboard** menu and select **Delete Current Dashboard**.

8.3 Sorting and Filtering Jobs

The job list on the Open Jobs, Ongoing Jobs, or Completed Jobs screens are displayed slightly different for the administrator. At the top of the screen, are two drop-down menus that let you view the list with jobs owned by a specific user and/or with jobs belonging to a specific category. The administrator can also choose to view all jobs/categories or jobs with no owner/category. The job list is refreshed according to the choices made in the drop-down menus. Each job is listed with its Job ID, Job Title, Category, Owner, Mail Type (plain or HTML), and "Authorization Due By" date.

Also, for the administrator the **Multiple Job Actions** feature is always enabled and can not be disabled.

Below the category selection box appears the **Advanced Filter Settings** link, which will take you to the Advanced Filter Settings screen that allows you to further define filters that are applied for the job list below (in addition to the filtering over the selected category).

If advanced filtering is disabled, then the control says "Filters are inactive" and the job list is not filtered further (except for the category filter, if applicable).

If advanced filtering is enabled, then the control says "Filters are active: See detailed settings below". In this case, the currently defined filter will be applied and the job list

below will show only jobs that fulfill the filtering condition (and are in the correct selected category).

In addition, the job list will display an additional filter row (with light-yellow background, just below the table header) that displays each filter in short form, as a reminder. For each column in the table, if a filter has been defined for that column, then the filter is displayed in the filter field below that column's header. If no filter is defined for a column, then the filter field of that column will be empty.

8.4 Archiving Delivered/Completed Jobs

To save server space and shorten jobs listings within the Maestro User interface, administrators can archive delivered jobs and jobs that have been closed after a failed delivery. Archiving a delivered or failed job removes the job from the system and saves it in a single ZIP archive file stored in a special archive folder on the system. Archived jobs cannot be viewed because all their tracking events are deleted and they are removed from any report data sources. As a result, any existing reports referencing them in their data sources will not display correctly.

The default archive folder of a LISTSERV Maestro installation on Windows is located along a path similar to: `\Program Files\L-Soft\Application Server\lui\archive`. On UNIX/Linux, the default archive folder is `~/lui/archive`. Although archived jobs are saved as ZIP files, little space will be saved because the archive folder exists on the same server or disk as the application. To save disk space, they can be moved from the server or disk where LISTSERV Maestro is installed. This can be done two different ways.

The first way is to change the default archive folder in the Administration Hub to point to a folder that is located on a different disk. The disk could be another disk on the same server, a mapped network drive (Windows), or a mounted NFS drive (UNIX/Linux) available on another server. By setting a different default folder for saving archived jobs within the Administration Hub, the list of archived jobs displayed on the Archived Jobs screen remains intact. All archived jobs in the folder will display in this list and can be imported back into LISTSERV Maestro if necessary. To change the default archive folder, see Section 5 [Settings for the Maestro User Interface](#).

The second way to move archived jobs from the disk or server where LISTSERV Maestro is located is to do so manually. Open the default archive folder. All archive files are ZIP files and have the job ID in the file name. Select the files and move them to a secondary storage medium such as different disk, a tape, a CD-ROM or similar. Once the file has been removed from the default archive folder it will not appear in the list of archived files. Files removed from the archive folder can be moved back in it at any time, and then will appear in the list of archived files. Once listed, the files will be available to import back into LISTSERV Maestro.

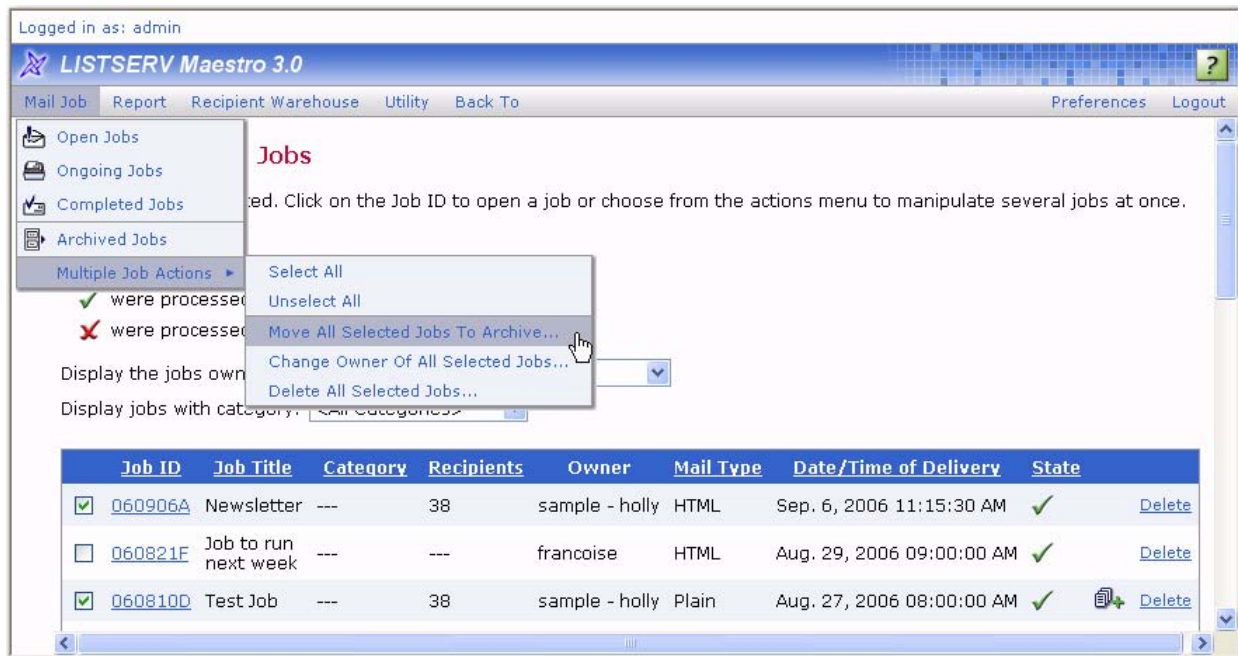
There are two ways to archive a delivered or failed job. Click on **Mail Job > Completed Jobs** from the Toolbar for administrative user account. The Completed Jobs screen opens listing all the previously delivered jobs. Then, either:

- Click on the Job ID link to select the job. The Job Administration screen opens. Click on the Job Archive tab, and then click the **[Move to Archive]** button. Click the **[Cancel]** button to cancel the operation and return to the Completed Jobs screen.

– OR –

- Check the box next to each job ID that you want to archive, and then click **Mail Job > Multiple Job Actions > Move All Selected Jobs to Archive**. The Move Selected Jobs to Archive screen opens. Click the **[Move to Archive]** button. Click the **[Cancel]** button to cancel the operation and return to the Completed Jobs screen. By using this method of archiving, you can archive more than one job at a time.

Figure 8-13 Archiving a Job from the Completed Jobs Screen



8.4.1 Auto-Archiving

LISTSERV Maestro now gives you the ability to automatically archive completed jobs that are older than a certain age. To define the default settings for auto-archiving, click on the **Global Component Settings** icon, then **Maestro User Interface**, and finally **Default Auto-Archive Settings**. The Auto-Archive Settings screen opens.

From this screen, define whether or not jobs will be automatically archived when the completed job reaches the “auto-archive age”. In the **Auto-Archive Age** field, enter the number of days old the completed job must be before it is automatically archived. If the age is set to 0, then the auto-archive setting is off. The age defined here will be the default for all accounts and groups. To define these settings for a specific user, see Section 7.3 [Managing User Rights](#) or Section 7.4.2 [Editing Component Specific Settings for Single and Group Users](#).

8.5 Importing Archived Jobs

Imported archived jobs are in a “frozen” state. The status and the contents of the job will not change from the moment it was placed in the archive. Any tracking events that arrive after the moment the job is archived will be discarded, even if the job is later imported.

To restore an archived job to the system, click **Mail Jobs > Archived Jobs**. The Archived Jobs screen opens with a listing of all the jobs currently present in the archive. From this screen, you can:

- Click on the **Job ID** link to select the job to be imported. Click **[OK]** to import the selected archived job. The Job Administration screen opens. Use the drop-down menu to select an owner for the imported archived job, and then click **[OK]**. Once restored, an imported job will be listed again in the Completed Jobs screen with its original Date and Time of Delivery (not the archived date). The job owner (and other group members, if applicable) can use imported archived jobs in tracking reports.

– OR –

- Check the box next to all the jobs you want to import, and then click **Mail Job > Multiple Job Actions > Import All Selected Jobs**. Click **[OK]** to import the selected archived jobs. By using this method of importing, you can import more than one job at a time. Imported jobs will have to be assigned a new Job Owner. From the Completed Jobs screen, display jobs with <No Owner>. Check the box next to the job(s) you want to assign a new owner to, and then click **Mail Job > Multiple Job Actions > Change Owner of All Selected Jobs**. From the Job Administration screen, use the drop-down menu to select an owner, and then click **[OK]**. Once restored, an imported job will be listed again in Completed Jobs screen with its original Date and Time of Delivery (not the archived date). The job owner (and other group members if applicable) can use imported archived jobs in tracking reports.

Figure 8-14 Importing a Job from the Archived Jobs Screen



8.6 Changing Job and Report Ownership

The user that initially creates a new job or report is the owner of that job or report. The owner is the only user with privileges (rights) to execute the following job or report related actions:

- Assign collaboration rights on the job to other group members

- Change the job information (job title and job ID-prefix)
- Delete the job
- Re-open, retry, or close a failed job.
- Assign collaboration rights on the report to other group members
- Delete the report

No other user can be granted owner rights within the same email job. Therefore, it is important that there is an owner for each job and each report because only the owner can execute these actions.

Under normal conditions, there will always be an owner because the initial creator will automatically become the owner. However, under certain circumstances, a job may lose its owner:

- If an account is deleted in the Administration Hub (and therefore also in the Maestro User Interface), and the **Keep data owned by the account** option was chosen, then all jobs and reports that were owned by that account will be without an owner.
- If an account that was a member of one group is changed so that it becomes a member of another group or not a member of any group at all, then all jobs and reports that were owned by that account will be without an owner (“orphaned”).

To reassign a job, go to the appropriate screen that would list that job. Click **Mail Job > Open Jobs** to open the Open Jobs screen for a listing of all the current jobs that have not been authorized for delivery yet; click **Mail Job > Ongoing Jobs** to open the Ongoing Jobs screen for a listing all the jobs that have been authorized for delivery; click **Mail Job > Completed Jobs** to open the Completed Jobs screen for a listing all the jobs that have already been delivered.

To reassign a report, click **Report > Reports** to open the Tracking Reports screen for a listing of all the currently defined reports.

If you want to view a job owned by a specific user, click the **Display the jobs owned by** drop-down menu and select that particular owner. Then, you can:

- Click on the **Job ID** link to select the job. The Job Administration screen opens. Click on the Job Owner tab (if necessary). Select a new owner for the job from the drop-down menu. Click **[OK]** to save the change, or **[Cancel]** to discard the change and return to the previous screen.

– OR –

- Check the box next to all of the jobs you want to reassign, and then click **Mail Job > Multiple Job Actions > Change Owner of All Selected Jobs**. Select **Change owner of all selected jobs**, and then use the drop-down menu to select an owner for the job(s). Click **[OK]** to save the change, or **[Cancel]** to discard the change and return to the previous screen.

Figure 8-15 Change Job Owner from Job Owner Tab

Job Administration

Job Owner | Job Archive

Change Job Owner

Select a new owner from the drop-down menu. Click [OK] to change the job owner.

Current job owner

Username: holly

Select a new owner: test1 [jht]

OK Cancel

If you've clicked **Report > Reports**, then the Tracking Reports screen opens with a listing of all the currently defined reports. If you want to view a report owned by a specific user, click the drop-down menu and select that particular owner. Then, you can:

- Click on the report title link to select the report. The Change Report Owner screen opens. Select a new owner for the report from the drop-down menu. Click **[OK]** to save the change, or **[Cancel]** to discard the change and return to the previous screen.

– OR –

- Check the box next to all of the reports you want to reassign, and then click the **Actions** link. Select **Change owner of all selected reports**, and then use the drop-down menu to select an owner for the report(s). Click **[OK]** to save the change, or **[Cancel]** to discard the change and return to the previous screen.

To delete more than one report or job at a time, simply check the box next to all of the reports/jobs that you want to delete, and then click the **Actions** link. Select **Delete all selected reports/jobs**, and then click **[OK]**. Click **[Cancel]** to discard the deletion and return to the previous screen.

Figure 8-16 Change Report Owner Screen

Change Report Owner

Report: Test Report

Select a user, and then click [OK] to change the report owner.

Current report owner

Username: holly
Group: sample

Select a new owner: wateam1 [sample]

OK Cancel

When an account is changed or ownership of a job or report is transferred, job and report rights are affected.

- **Account name is changed; account not in a group** – The account retains ownership of all jobs and reports. The account remains the owner of all previous jobs and reports. Since the account is not a member of a group, neither before nor after the change, team collaboration rights on the jobs and reports are not affected.
- **Account name is changed; account member of a group** – The account retains ownership of all jobs and reports and the jobs and reports remain in the same group. The account retains ownership of all previous jobs and reports. The team collaboration rights on the jobs and reports are not affected. All members in the same group that had any job or report privileges before the change have the same privileges after the change.
- **Group is added to a non-group account** – The account retains ownership of all jobs and reports and transfers them to the new group. The account remains owner of all previous jobs and reports. Since the account was not a member of a group before the change, no team collaboration rights on the jobs and reports are affected. However, after the change, since the account is now a member of a group, the owner of the account can assign team collaboration rights on any of the jobs or reports to any of the other members in the group.
- **Group is removed from a group account** – The account loses ownership of all jobs and reports. Jobs and reports remain in the old group. All previous account privileges for any jobs or reports are removed. This means that the account loses all team collaboration rights on all jobs and reports that are owned by other members of the account's old group. The account also loses ownership of all jobs and reports that the user of the account created while the account was still a member of the old group. The team collaboration rights of other members of the old group are unaffected. However, these jobs no longer have an owner, since the old owner left the group. The administrator should set a new owner at this point.
- **Group account becomes member of different group** – The account loses ownership of all jobs and reports. Jobs and reports remain in the old group. All previous account privileges for any jobs or reports are removed. This means that the account loses all team collaboration rights on all jobs and reports that are owned by other members of the account's old group. The account also loses ownership of all jobs and reports that the user of the account created while the account was still a member of the old group. As a result, the account joins the new group as a “fresh” member, without any team collaboration or job or report ownership rights. The jobs and reports that were created by the user of the account while still in the old group remain in the old group. The team collaboration rights that other members of the old group may have on those jobs and reports are unaffected. However, these jobs no longer have an owner, since the old owner left the group. The administrator needs to set a new owner at this point.
- **Ownership of a job or report is transferred; previous owner not in a group** – The new account acquires all ownership rights on the job or report. The original owner loses all rights, including ownership.
- **Ownership of a job or report is transferred; previous owner in a different group from new owner**
The new account acquires all ownership rights on the job or report. For all other

accounts (the previous owner and the members of the old group), ownership or team collaboration rights are removed.

- **Ownership of a job or report is transferred; previous owner in same group as new owner**

The new account acquires all ownership rights on the job or report. The original owner loses all rights, including ownership. However, any other accounts that may have team collaboration rights on the job or report do retain these rights – they are preserved.

8.7 Changing Sender Profile, Drop-In, and Content Template Ownership

Drop-ins, content templates, and sender profiles are utility items in LISTSERV Maestro, for which ownership is handled as described below.

If such an item is created by a user that is not a member of a group, then this item is owned by that user, and can be used only by him.

If the item is created by a user that is a member of a group, then this item is owned by the group, i.e. all accounts in the group can use this item, and all accounts with the necessary right to create items of that type can also modify or delete it.

An item can also lose its owner and become ownerless, if one of the following happens:

- If an item is owned by a non-group account, and that account is deleted with the **Keep data owned by the account** option selected.
- If an item is owned by a group and the group is deleted with the **Keep data owned by the group and its accounts** option selected, or the last account of that group is deleted with the **Keep data owned by the account or group** option selected, or the last account of that group is moved out of the group (into another group or made into an account without a group).

The ownership of any item, whether it already has an owner or not, can be changed by the administrator using the special Administrator User Account. The administrator can also delete any item.

To change ownership for or delete a drop-in, select **Utility > Drop-Ins**. The Manage Drop-In Content Elements screen opens.

To change ownership for or delete a content template, select **Utility > Content Templates**. The Manage Content Templates screen opens.

To change ownership for or delete a sender profile, select **Utility > Sender Profiles**. The Manage Sender Profiles screen opens.

At the top of the screen, there is a drop-down menu with all available owners, e.g. all available non-group accounts and all available groups, as well as a special “<No Owner>” entry. From this drop-down menu, select the owner you’d like to view.

Below the drop-down menu appears a table with all items that are currently owned by the selected owner account or group.

Click on an item to change the ownership or delete. The Change Owner screen opens. Select a new owner from the drop-down menu. Click **[OK]** to save the change, or **[Cancel]** to discard the change and return to the previous screen. To delete, click the **[Delete]** button.

Figure 8-17 Example: Changing the Owner of a Sender Profile



If an item is moved to a new owner, and the new owner already has an item with the same name, then the moved item is automatically renamed to give it a unique name. For example, if an item named “sample” is moved to an owner who already has an item with that name, the moved item is renamed “sample1”. If “sample1” is also in use, the moved item will be renamed “sample2”, and so on until a unique name is created.

Ownership of an item is affected when an account is changed in the following ways:

- **Account name is changed, account not in a group**
The account keeps ownership of all items of which it previously was the owner.
- **Account name is changed, account member of a group**
Since the items are owned by the account’s group anyway, ownership is not affected, meaning that they are still owned by the same group as before.
- **Group is added to a non-group account**
Ownership of all items previously owned by the account is transferred to the group that the account joins. All members in that group will then share ownership of these items. The account also gains ownership of all items that previously existed in the group it joins.
- **Group is removed from a group account**
Since the group owns the items, the account loses access to them. They stay in the old group. The account that no longer belongs to a group does not have ownership of any items until new ones are created with this account.
- **Group of a group account is changed – account becomes member of different group**
Since the old group owns the items, the account loses access to them. They stay in the old group. The account gains ownership of all items that previously existed in the new group it joins.

8.8 Changing Ownership of the Entire Recipient Warehouse

A recipient warehouse is the container of all datasets, target groups, hosted lists, and lookup tables of a certain owner. There are two possible types of owners:

- **Recipient warehouse owned by a user which is not a member of a group:** Each user which is not a member of a group owns his individual recipient warehouse, and thus all datasets, hosted lists and lookup tables in it. Only this owner can use the objects in his recipient warehouse.
- **Recipient warehouse owned by a group:** Each group owns its group's recipient warehouse with all its datasets, hosted lists and lookup tables. All member in the group (provided they own the necessary user right) may use the objects in this recipient warehouse. Any dataset, hosted list or lookup table which is created by any user in the group will be part of this one group-owned recipient warehouse.

A recipient warehouse (and the objects in it) can lose its owner and become ownerless, if one of the following happens:

- If a recipient warehouse is owned by a non-group account, and that account is deleted with the **Keep data owned by the account** option selected.
- If a recipient warehouse is owned by a group and the group is deleted with the **Keep data owned by the group and its accounts** option selected, or the last account of that group is deleted with the **Keep data owned by the account or group** option selected, or the last account of that group is moved out of the group (into another group or made into an account without a group).

The ownership of a recipient warehouse, no matter if it still has an owner or not, can be manipulated by the administrator. This means that the administrator can assign a recipient warehouse to a new owner (even if the recipient warehouse still has an owner). However, the administrator can only assign the whole recipient warehouse to a new owner, not the individual datasets, target groups, hosted lists, or lookup tables in it. In reality, the recipient warehouse is not truly assigned to a new owner, but simply the objects in it are moved from the recipient warehouse of the old owner to the recipient warehouse of the new owner. This means that the old owner (if still existing) remains the owner of their own recipient warehouse (only it is now empty), and the moved objects are simply added to the already existing recipient warehouse of the new owner.

The administrator can also delete a recipient warehouse (of a certain owner) or any object in it. Recipient warehouse objects can be deleted on several levels:

- **Delete the whole recipient warehouse:** Actually, this does not really delete the recipient warehouse, but it simply deletes all datasets, target groups, hosted lists, and lookup tables in it, i.e. the recipient warehouse remains, but is now empty.
- **Delete a dataset in the warehouse:** This deletes the dataset and hosted lists.
- **Delete a recipient target group:** This deletes the target group.
- **Delete a hosted list:** This deletes only the selected hosted list.
- **Delete a lookup table:** This deletes only the selected lookup table. However, a lookup table can only be deleted in this fashion if it is not currently used by any dataset or hosted list.

If the objects in a recipient warehouse are moved to a new owner, and the recipient warehouse of the new owner already contains objects, and an object with the same name as one of the moved objects already exists in the recipient warehouse of the new owner,

then the moved object is renamed to give it a unique name. For example, an object called “sample” is renamed to “sample1”, unless this name is also already in use, in which case it is renamed to “sample2”, etc.

The following list describes how ownership of the objects in a recipient warehouse is affected when an account is changed:

- **Account name is changed, account not in a group**
The account keeps ownership of the recipient warehouse and all objects in it.
- **Account name is changed, account member of a group**
Since the recipient warehouse is owned by the account’s group anyway, its ownership is not affected, i.e. it is still owned by the same group as before.
- **Group is added to a non-group account**
All objects in the recipient warehouse previously owned by the account are transferred to the recipient warehouse of the group that the account joins. Thus, all members in that group will then share ownership of these new objects in their recipient warehouse. The account also gains access to the group recipient warehouse and all objects in it (i.e. both the ones that were already there and the ones that moved with him). The previous (now empty) recipient warehouse of the account is discarded as it is no longer needed.
- **Group is removed from a group account**
Since the recipient warehouse is owned by the group, the account loses access to it. Instead, the account gets his own new recipient warehouse only that it is initially empty. Thus, the account which now no longer belongs to the group does not have ownership of any objects, until new ones are created with this account.
- **Group of a group account is changed – account becomes member of different group**
Since the recipient warehouse is owned by the old group, the account loses access to it and the objects in it. It stays with the old group. However, the account gains access to the objects in the recipient warehouse owned by the new group it joins.

8.8.1 Changing Target Group Ownership

As with sender profiles, content templates, and drop-ins, target groups may have either single user or group ownership. Single user target groups are owned by an individual user, and may not be used by anyone else. If a group member creates items, they may be used and modified by other users within the group who have the necessary rights.

An item can lose its owner and become “owner-less” under certain circumstances:

- If an item is owned by a non-group account, and that account is deleted with the **Keep data owned by the account** option selected.
- If an item is owned by a group and the group is deleted with the **Keep data owned by the group and its accounts** option selected, or the last account of that group is deleted with the **Keep data owned by the account or group** option selected, or the last account of that group is moved out of the group (into another group or made into an account without a group).

The ownership of any item, whether it already has an owner or not, can be changed by the administrator using the special Administrator User Account. The administrator can also delete any item.

To change the ownership of a target group category or an individual target group, click **Recipient Warehouse > Target Groups**. The Change Target Group Owner screen opens.

Figure 8-18 Change Target Group Owner Screen

Target Group Category	Owner
<No Category>	Account francoise
<Hosted Lists Default Target Groups>	Account francoise
<Hosted Lists Default Target Groups>	Group sample
Target Groups	Group sample

Click the name of the target group category to select it, and then select a new owner from the drop-down menu.

Figure 8-19 Change Ownership of a Target Group Category

Target Group	Description
Target Group 1	Target Groups

Click **[OK]** to save the change, or **[Cancel]** to discard the change and return to the previous screen.



Notes: To delete a target group category, click the **[Delete]** button. If you delete the category, then all target groups within that category will also be deleted. If you change the owner of a category, then all target groups within that category will automatically change to the new owner as well. The <No Category> target group category cannot be deleted or moved to a different owner because it exists in the scope of all owners.

To change the ownership of an individual target group, click the target group name in the list at the bottom of the Change Target Group Owner screen. The next Change Target Group Owner screen opens. From this screen, select the new owner from the drop-down menu. Click **[OK]** to save the change, or **[Cancel]** to discard the change and return to the previous screen.

The target group will be moved to the new owner. If the target group was within a category, and a category with the same name already exists in the scope of the new owner, then the target group will appear in it. If the new owner does not have a category with the same name, then one will be automatically created so that the target group can be moved into it. If the new owner already has a target group with the same name within the category, the new target group will automatically be renamed. The following renaming scheme will be used: If the original name was "sample", and a target group with this name already exists, then the new name will be "sample1". If this name also is in use, "sample2" will be used, and so on.

Figure 8-20 Change Ownership of an Individual Target Group

Change Target Group Owner

Selected Target Group: D-HLL All Subs <No Category>

All active subscribers to D-HLL (all topics)

Select a new account or group from the drop-down menu. Click OK to make that account/group the new owner of the target group.

Current target group owner

Account: francoise

Select a new owner: Group test

OK Cancel Delete



Note: To delete an individual target group, click the **[Delete]** button.

8.8.2 Changing Ownership of Recipient Datasets and Lookup Tables

To change the ownership of a Recipient Dataset, click **Recipient Warehouse > Datasets & Lookup Tables**. The Recipient Datasets and Lookup Table Administration screen opens.

Select a new owner from the drop-down menu. Click **[OK]** to save the change, or **[Cancel]** to discard the change and return to the previous screen.



Note: To delete a Recipient Dataset, click the **[Delete]** button.

Figure 8-21 Change Ownership of a Recipient Dataset

Recipient Dataset and Lookup Table Administration

Select a new account or group from the drop down menu.
Click OK to make that account/group the new owner of the datasets and lookup tables.

Current Owner: Group *sample*

Select a New Owner:

Datasets Owned: 1

Lookup Tables Owned: 0

Below you see a list of all datasets owned.
Click on a dataset to select it.

Recipient Datasets

[Sample dataset](#)

The Recipient Datasets and Lookup Table Administration screen also lists all datasets and lookup tables of the current owner.

Lookup tables are listed with their name and a short list of references (how many fields of how many datasets or mailing lists use the lookup table). For lookup tables without any references, the name is clickable. Click to open the Lookup Table Administration screen for that lookup table (for lookup tables that have at least one reference, no further administration is possible; therefore, their names are not clickable). From the administration screen, view the name of the Lookup Table, a description of the table, and the current owner of the table. From here, it is also possible to **[Delete]** the Lookup Table. Click **[Cancel]** to return to the previous screen.

Datasets are listed with their name as a link. Click the link to open the administration screen for that dataset.

Figure 8-22 Recipient Datasets Administration Screen

Recipient Dataset Administration

Click the Delete button to remove the selected dataset.

Recipient Dataset: *My Test Dataset*

Description: Test

Current Owner: Account *francoise*

Hosted Lists in Dataset: 5

Below you see a list of all hosted lists in the dataset.
Click on a list to select it.

Hosted Lists

[A-HLL](#)

[D-HLL](#)

[HRL](#)

[HRL2](#)

[U-HLL](#)



Note: It is not possible to re-assign individual datasets or lookup tables. When you choose to re-assign a specific owner's datasets and lookup tables, then you are re-assigning that owner's entire Recipient Warehouse.

From the Recipient Dataset Administration screen, it is possible to **[Delete]** the Lookup Table. Click **[Cancel]** to return to the previous screen.

If the dataset contains Hosted Lists, click on the individual list names to open the Hosted Lists Administration screen. From that screen, view the name of the Hosted List, a description of the list, the current list owner, and the Recipient Dataset to which the Hosted List belongs. From here, it is possible to **[Delete]** the Hosted List from the dataset or **[Cancel]** to return to the previous screen.

Figure 8-23 The Hosted List Administration Screen

Hosted List Administration

Click the Delete button to remove the selected list from the recipient dataset.

Hosted List: [A-HLL](#)

Description: A test HLL

Current Owner: Account [francoise](#)

Recipient Dataset: [My Test Dataset](#)

Section 9 Defining External Database Connections

The Maestro User Interface component of LISTSERV Maestro uses a “system database” to store its working data, outlined in Section 10 [The System Database](#). This may be either the default “internal” MySQL database, or some external database.

The Maestro User Interface can also be configured to access an external “user database” to retrieve existing information to build recipient lists in the recipients wizard or target groups wizard, or to create drop-in content elements.

Multiple databases managed by the same or different DBMS software can be configured as user databases so that recipient data and drop-in content elements can be accessed from many sources. The user database(s) may be on the same database server as the system database, or on different ones. By configuring LISTSERV Maestro to be able to access different databases, institutional data can be retrieved from different sources, allowing for great flexibility.

The following DBMS products have been tested and are compatible with LISTSERV Maestro:

- Microsoft® SQL Server 7.0 and 2000
- Oracle® 8i Enterprise/Standard Edition Release 3 (8.1.7)
- Oracle® 9i Release 2 (9.2.0.3) & (9.2.0.1)
- Oracle® 10g Enterprise/Standard Edition Release 1 (10.1.0.2.0)
- DB2® Universal Database V7.2 and V8.2
- MySQL® 4.1.7, MySQL 4.0.22, MySQL 3.23.42
- Any ODBC compliant database can be used for read-only purposes to retrieve recipient lists and drop-in content elements.

DBMS versions comparable to the versions listed above should also be compatible with LISTSERV Maestro (this is particularly true for versions later than those listed). Versions earlier than those listed above are not officially supported.

LISTSERV Maestro communicates with external databases with so called “*Plugins*” and drivers. If an external database is going to be used for the system database or the user database, the appropriate driver must be installed and the plugin must be configured first.

Before an external database can be invoked, either as the system database in the HUB System Database Connection screen, or as a user database in the LUI Data Warehouse, Recipient Definition, Target Group Definition, or Drop-in Definition screens, LISTSERV Maestro must know how to access the particular DBMS software managing the database in question.

The following steps need to be taken once for each DBMS package, which will make any databases running under that software available:

- Install the driver for the database on the server where the Maestro User Interface (LUI) is installed. See Section 9.1 [Available Database Plugins](#) for details.
- Register the appropriate plugin in the Administration Hub (HUB) component. See Section 9.2 [Registering a Database Plugin](#) for more information.



Important: Connection details for user databases are defined in the recipients target groups wizard, or the recipients wizard in the Maestro User Interface during the recipient definition of a job. Do not enter connection details in the HUB for user databases. Connection details are entered in the HUB only for the external system database. After a restart, any database connection details entered in the Global Components Settings will change the system database.

9.1 Available Database Plugins

The Maestro User Interface is a Java server application that uses JDBC to connect to the configured database. Therefore, it is usually necessary to install a compatible JDBC driver for the database. Each database plugin (see Section 9.2 [Registering a Database Plugin](#) for more information) has been developed to use exactly one JDBC driver. There may be several plugins for the same DBMS, each of which uses a different driver to access that DBMS. The specific plugin to be used depends on the DBMS and the JDBC driver available for that DBMS.



Important: After installing a new JDBC driver into LISTSERV Maestro (see descriptions below), it is necessary to restart LISTSERV Maestro to make it aware of the new driver. On Windows, you must also re-install the LISTSERV Maestro Windows service. To do so, execute the following command after installing the driver:

```
[maestro_install_folder]/commands/InstallService.cmd
```



Note: The plugins available at the time this document was written support nine different drivers for four different databases as well as the ODBC-driver (as a read-only plugin only) that in turn allows access to any database or other data source that has an ODBC driver available.

9.1.1 The IBM DB2 V8.2 Thin Driver Database Plugin

This plugin is used for connecting to the DB2 V8.2 database and uses the IBM DB2 V8.2 thin driver.

- **Plugin class name:** `com.lsoft.lui.db.ibm.DB2V82ThinDriverPlugin`
- **How to install the driver:** The driver comes in form of four files which are found in the installation folder of the DB2 V8.2 database (license conditions from IBM may apply):

```
[db2_install_folder]/SQLLIB/java/db2jcc.jar
[db2_install_folder]/SQLLIB/java/db2jcc_javax.jar
[db2_install_folder]/SQLLIB/java/db2jcc_license_cu.jar
[db2_install_folder]/SQLLIB/java/db2policy.jar
```

Simply copy these files into the “lib” folder in the LISTSERV Maestro installation:

```
[maestro_install_folder]/lib
```


9.1.2 The IBM DB2 V7.2 Native Driver Database Plugin

This plugin is used for connecting to the DB2 V7.2 or V8.2 database and uses the IBM DB2 V7.2 or V8.2 native driver. (For accessing DB2 V8.2, it is recommended to use the thin driver instead, see above.)

- **Plugin class name:** `com.lsoft.lui.db.ibm.DB2V72DriverPlugin`
- **How to install the driver:** The driver comes as part of the DB2 V7.2 or V8.2 database installation (license conditions from IBM may apply):

To install the driver, first install the DB2 run-time clients (from the runtime folder of the installation package) on the server where you want to run the Maestro User Interface. Use the client to connect to a database on the DB2 server (see Section 10.2.4 [Preparing DB2 as the System Database](#)). It is important that you install the client on the Maestro User Interface server, *not* on the server where the database is installed (except of course, if both components happen to be on the same server).

Only for DB2 V7.2:

On Windows, stop the JDBC DB2 Applet Server service.

Then run the batch command file “`usejdbc2.bat`” (Windows) or “`usejdbc2.sh`” (Linux) from the “`java12`” subfolder in the DB2 run-time client installation, i.e. execute the file:

```
[ibm_install_folder]/SQLLIB/java12/usejdbc2.bat | sh
```

This script prepares the DB2 runtime environment for JDK1.2 and later, which is required for a fully functional LISTSERV Maestro installation.



Note: The batch command file creates a different version of the JDBC driver file with the name “`db2java.zip`”. Make sure to use the newly created file when proceeding to the next step.

On Windows, start the JDBC DB2 Applet Server service.

Both for DB2 V7.2 or DB2 V8.2:

Copy the file “`db2java.zip`” from the run-time client installation to the LISTSERV Maestro installation:

Copy the file “`db2java.zip`” from the “`java`” folder in the DB2 run-time client installation:

```
[ibm_install_folder]/SQLLIB/java
```

into the “`lib`” folder in the LISTSERV Maestro installation:

```
[maestro_install_folder]/lib
```

9.1.3 The MySQL ConnectorJ Driver Database Plugin

This plugin is used for connecting to the MySQL database of version 4.1 and later (tested until 5.0, at the time this was written). This plugin uses the ConnectorJ MySQL driver, which is installed together with LISTSERV Maestro.



Note: This plugin can not be used to connect to MySQL instances with a version earlier than 4.1. If you want to connect to such older MySQL versions, you need to use the MySQL L-Soft Driver Database Plugin (see below).

- **Plugin class name:**
`com.lsoft.lui.db.mysql.MySQLConnectorJDriverPlugin`
- **How to install the driver:** The driver is pre-installed together with LISTSERV Maestro.

9.1.4 The MySQL L-Soft Driver Database Plugin

This plugin is used for connecting to the MySQL database of version 3.23.42 (or later 3.23.x builds) or version 4.x. This plugin uses the L-Soft MySQL driver, which is installed together with LISTSERV Maestro.



Note: This plugin can not be used to connect to a MySQL instance with a version of 4.1 or later. If you want to connect to such newer MySQL versions, you need to use the MySQL ConnectorJ Driver Database Plugin (see above).

- **Plugin class name:** `com.lsoft.lui.db.mysql.MySQLDriverPlugin`
- **How to install the driver:** The driver is pre-installed together with LISTSERV Maestro.

9.1.5 The Oracle 8i, 9i, and 10g Thin Driver Database Plugin

These plugins are used for connecting to the Oracle database of version 8*i*, 9*i*, or 10g. These plugins use the Oracle Thin driver.



Note: The newer driver for 10g may also work for 9i and 8i and is usually more efficient. It is therefore recommended that you use the newer 10g driver even when connecting to Oracle 9i or 8i.

- **Plugin class name:**
`com.lsoft.lui.db.oracle.Oracle8iThinDriverPlugin`
- **How to install the driver for Oracle 8i:**

The driver comes in form of a file called “classes12.zip”. Simply copy this file into the “lib” folder in the LISTSERV Maestro installation:

```
[maestro_install_folder]/lib
```

The driver can be downloaded from the Oracle Technology Network. Look for the releases for Oracle 8*i* and the classes12.zip file. License conditions from Oracle may apply.

- **How to install the driver for Oracle 9i or 10g:**

The driver comes in form of a file called “ojdbc14.jar”. Simply copy this file into the “lib” folder in the LISTSERV Maestro installation:

```
[maestro_install_folder]/lib
```

The driver can be downloaded from the Oracle Technology Network. Look for the releases for Oracle 9i (or 10g) and the `ojdbc14.jar` file. License conditions from Oracle may apply.

9.1.6 The SQL Server jTDS Driver Database Plugin

This plugin is used for connecting to the SQL Server database of version 6.5, 7.0, 2000 or 2005. This plugin uses the free open-source jTDS driver.

- **Plugin class name:** `com.lsoft.lui.db.sqlserver.JTDSDriverPlugin`
- **How to install the driver:** From the binary distribution download, copy the file “`jtds-1.2.jar`” into the “`lib`” folder in the LISTSERV Maestro installation:

```
[maestro_install_folder]/lib
```

(At the time this document was written, version “1.2” was the most current version. When a new version is released, the name of the `jar` file that needs to be copied into the `lib` folder will probably change accordingly.)

The driver can be downloaded from the jTDS website: <http://jtds.sourceforge.net>. License conditions may apply.

9.1.7 The SQL Server Microsoft Driver Database Plugin

This plugin is used for connecting to the SQL Server database of version 2000 or 2005. This plugin uses the SQL Server driver from Microsoft, which comes in two versions – the older driver for SQL Server 2000 only and the newer driver for SQL Server 2005 (which also works with SQL Server 2000). From the two drivers, the plugin will automatically choose the driver which is currently installed on your system. If both drivers are installed, the plugin will choose the newer driver, which works both with SQL Server 2005 and 2000.

- **Plugin class name:** `com.lsoft.lui.db.sqlserver.MSSQLDriverPlugin`
- **How to install the driver for SQL Server 2000:** This is the older JDBC driver from Microsoft which works with SQL Server 2000 only. The driver comes in form of a Windows install file. Execute the install file on any computer you like. What is important for the Maestro User Interface is not that the driver is installed on the same computer, but that you copy the following files from the installation folder of the driver to the installation folder of LISTSERV Maestro (shown for a default installation of the driver):

Copy the files “`msbase.jar`”, “`mssqlserver.jar`” and “`msutil.jar`” from the “`lib`” folder in the SQL Server JDBC-driver installation:

```
\Program Files\Microsoft SQL Server 2000 driver for JDBC\lib
```

into the “`lib`” folder in the LISTSERV Maestro installation:

```
[maestro_install_folder]/lib
```

The driver can be downloaded from Microsoft’s SQL Server 2000 website. License conditions from Microsoft may apply.

- **How to install the driver for SQL Server 2005:** This is the newer JDBC driver from Microsoft which works with SQL Server 2005 and also SQL Server 2000.

The driver comes in the form of a self-extracting * .exe file. Extract the file to a suitable temporary location. Among the extracted files, you will find one file called “sqljdbc.jar”. Copy this file into the “lib” folder in the LISTSERV Maestro installation:

```
[maestro_install_folder]/lib
```

The driver can be downloaded from Microsoft’s SQL Server 2005 website. License conditions from Microsoft may apply.

9.1.8 The SQL Server i-net SPRINTA Driver Database Plugin

This plugin is used for connecting to the SQL Server database of version 6.5, 7.0, 2000, or 2005. This plugin uses the SPRINTA SQL Server driver from i-net software.

- **Plugin class name:** `com.lsoft.lui.db.sqlserver.SPRINTADriverPlugin`
- **How to install the driver:** From the SPRINTA download/installation, copy the file “Sprinta2000.jar” (older driver versions) or “Sprinta.jar” (newer driver versions) into the “lib” folder in the LISTSERV Maestro installation:

```
[maestro_install_folder]/lib
```

The driver can be purchased and downloaded from i-net software: <http://www.inetsoftware.de>. License conditions from i-net software may apply.



Note: The evaluation version of this driver which is (or was at the time this was written) available for download, contains a limitation of the number of concurrent database connections that will make the Maestro User Interface fail during operation. The evaluation version is therefore not supported for use with the Maestro User Interface.

9.1.9 The ODBC Driver Database Plugin

This plugin is used for connecting to any ODBC compliant database or data source, and uses the ODBC driver which is part of Java, which is installed together with LISTSERV Maestro.

- **Plugin class name:** `com.lsoft.lui.db.odbc.ODBCDriverPlugin`
- **How to install the driver:** The driver is pre-installed together with LISTSERV Maestro.

The ODBC-driver plugin is a *read-only plugin*. As such, it can only be used to read recipient data or drop-in content data. It cannot be used for the system database connection, or to create Hosted Recipient Lists. After registering this plugin it will not appear in the list of available drivers on the system connection page in the Administration Hub. However, it will appear in the corresponding lists of the recipient wizard, target group wizard, and database drop-in page.

On Windows installations, this driver is automatically installed together with LISTSERV Maestro, so the only step required to make this plugin available for usage is to register it as described in Section 9.2 [Registering a Database Plugin](#).

The ODBC driver plugin operates differently when compared to the other database plugins. The other plugins bind a specific JDBC driver to LISTSERV Maestro, allowing access to the specific database for which the JDBC driver has been written. Database access then goes through three layers, from the plugin into the JDBC driver and from there into the database as shown below.

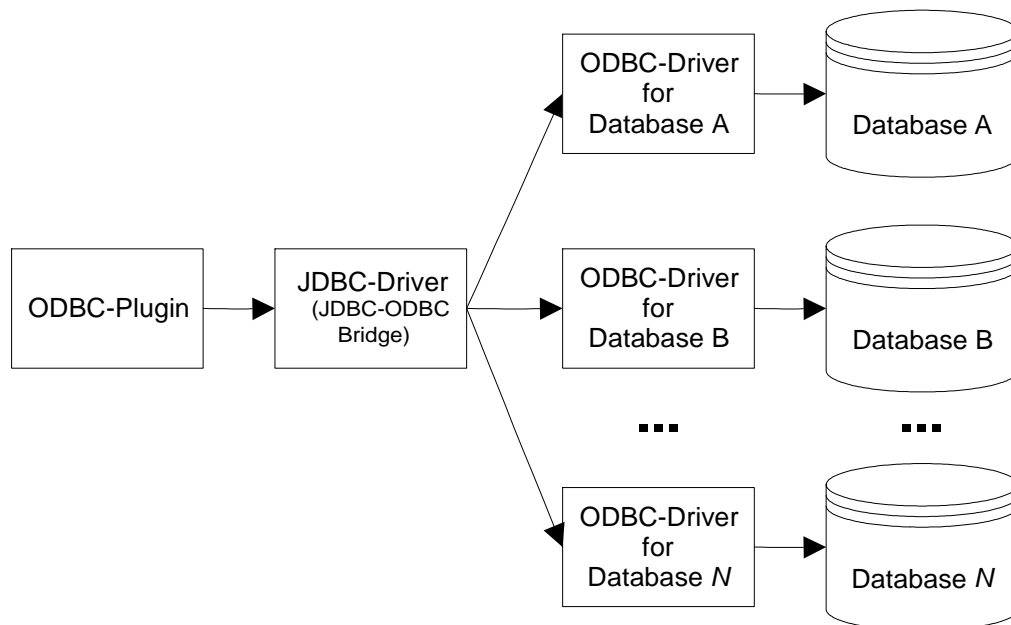
Figure 9-1 JDBC Driver Layers



In contrast, the ODBC driver plugin has one more layer shown below. This plugin binds the JDBC-ODBC bridge driver to LISTSERV Maestro, allowing access to any ODBC driver. The plugin by itself does not define which database is accessible through it. It does however, define that anything that has an ODBC driver available to it is accessible. The ODBC driver for the database in question must also be supplied in addition to the plugin. Installation of the ODBC driver depends on the system and the OS in use. Please see the appropriate documentation for the ODBC driver and the operating system.

Database access goes through four layers starting with the ODBC plugin to the JDBC-ODBC bridge, to the ODBC driver, and ending with the database.

Figure 9-2 ODBC Plugin Layers



The performance of LISTSERV Maestro when using this driver is directly dependant on the ODBC driver used for the database in question. Accessing a database through an ODBC driver that is programmed inefficiently will impact the performance of LISTSERV Maestro. For example, if the ODBC driver uses up a lot of memory when doing large

selects, LISTSERV Maestro may be subjected to a memory shortage caused by the ODBC driver. In that case, the driver is not usable unless it can be used to make smaller selects, or the server's memory is upgraded accordingly.



Important: Extensive testing with the ODBC driver(s) before employing in a production setting is recommended to determine the impact on memory and CPU usage.



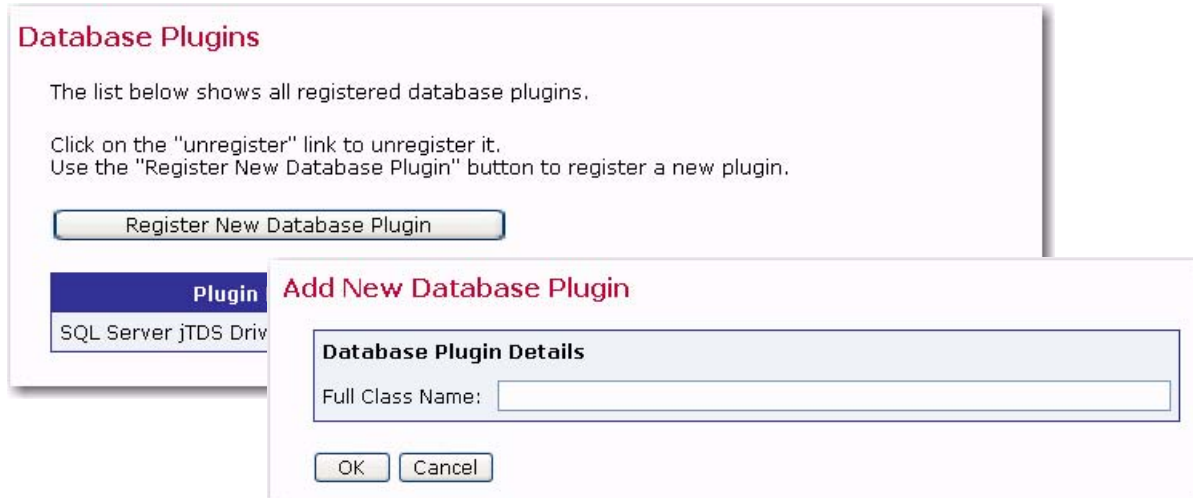
Tip: The term “database”, when used with ODBC, is interpreted quite broadly. ODBC drivers for data sources, such as plain text files or if Microsoft Excel files exist, turning them into “databases” in the sense that they can be used to create recipient lists and drop-in content. See the L-Soft White Paper entitled *Importing Data from Microsoft Excel into LISTSERV Maestro* for an example.

9.2 Registering a Database Plugin

LISTSERV Maestro uses “database plugins” to give access to different JDBC drivers (and through them to different databases) available to the Maestro User Interface. Before a plugin can be used, it must first be registered in the list of known plugins. Some plugins are already pre-registered when LISTSERV Maestro is installed, while others need to be registered after the corresponding JDBC driver has been installed.

To register a new plugin, log into the Administration Hub and click on the **Global Settings** menu, select **Maestro User Interface**, and then finally select **Database Plugins**. Click on the **[Register New Database Plugin]** button. In the text box, enter the full class name of the plugin to be registered (see Section 9.1 [Available Database Plugins](#)).

Figure 9-3 Adding a Database Plugin



Click the **[OK]** button to submit the class name. If the plugin was registered correctly, it will now appear in the list of plugins. If there was a problem during the registration, an error message describing the problem will appear. The most probable causes for problems are misspellings of the class name (which is case sensitive).

Section 10 The System Database

LISTSERV Maestro uses a “system” database to store its working data – recipient profiles, job ID numbers, tracking information, and so on. An “internal” database (based on MySQL) is included as part of the application and may be used as the system database. Using this internal database will allow the application to run “out-of-the-box”. An optional external database may be configured in place of the default internal database if desired. Switching the system database from internal to external and vice versa can be done at any time as long as both databases are accessible to LISTSERV Maestro during the transfer.



Notes: See <http://www.mysql.com> for details about MySQL 4.1.7 and its features to determine if it meets your requirements and expectations for a production database. Reviewing the features will help you decide if you want to use the internal database or if you prefer to install a different external system database.

The internal MySQL database is installed on the same server as the Maestro User Interface component (if installation of MySQL was not deselected during setup). You can find the database binaries, including the server and client tools, in this folder on that server: `[maestro_install_folder]/lui/database/bin`

LISTSERV Maestro is also supports various other databases, of various versions.

- IBM DB2 Universal Database
 - V7.2
 - V8.2
- Microsoft SQL Server
 - 6.5 (as user-database only)
 - 7.0 (recommended as user-database only, see note below)
 - 2000
 - 2005
- MySQL
 - 3.23.42 and later 3.23.x builds (recommended as user-database only, see note below)
 - 4.0 (recommended as user-database only, see note below)
 - 4.1
 - 5.0
 - 5.1
- Oracle
 - 8i
 - 9i
 - 10g

New

Any comparable/compatible versions should also be usable. Any versions earlier than the ones mentioned here are not officially supported as system databases.

You can install any of these databases (or a compatible version) on the server where you have installed the Maestro User Interface component (or on any other server that is reachable over the network) and configure the Maestro User Interface to use this database instead of the internal database. (Please refer to the documentation of the database in question for details about how to install and configure the database).



Note: SQL Server 7.0, MySQL 3.23.x and MySQL 4.0 as a system database are also supported but for backwards compatibility with existing installations only. These versions are not recommended to be used with new installations and it is strongly recommended to upgrade existing installations to a newer SQL Server or MySQL version (or switch to one of the other supported databases).

10.1 Configuring the External System Database

In order to use an external database as the system database, the database must be prepared for use with LISTSERV Maestro, and LISTSERV Maestro must be configured to use the prepared database. LISTSERV Maestro must transfer all data from the old database to the new database in order to switch from the internal system database to an external database, even before LISTSERV Maestro is first used. Both databases must be running and accessible to LISTSERV Maestro during the transfer.



Important: If one external database is being switched for another external database, both external databases must be accessible to LISTSERV Maestro simultaneously so that the data can be automatically transferred.

Follow these steps to configure an external database:

1. Install the DBMS and prepare the database to be used. Follow the manufacturer's instructions to install the DBMS. See Section 10.2 [Preparing the System Database](#) for details on how to create and prepare the system database. Some instructions are different depending on the DBMS in use.
2. If changing the system database for an installation that has been in use, make a backup before making any changes so that the current data may be recovered in case of errors. See Section 11 [Saving and Restoring a Backup](#) for more information.
3. Install the corresponding JDBC driver on the server where the Maestro User Interface (LUI) is installed. See Section 9.1 [Available Database Plugins](#) for details.
4. Register the appropriate "Plugin" in the Administration Hub (HUB) component. See Section 9.2 [Registering a Database Plugin](#) for more information.
5. Define the System Database Connection as follows:
 - a. Log into the Administration Hub. Click on the **Global Settings** menu, select **Maestro User Interface**, and then finally select **System Database Connection**. The System Database Connection screen opens.
 - b. Select **The following external database is used as the System Database**.

- c. Click the **Database Plugin** drop-down menu and select an external database to use as the System Database.

Figure 10-1 System Database Connection Screen

- d. Fill out the database connection detail fields. These fields are different for each type of database selected.

Figure 10-2 System Database Connection Details Screen for DB2

- e. Click **[OK]** to submit the change.
6. Shut down LISTSERV Maestro and then restart.

7. Wait for the transfer of data between the two databases to be completed. Depending on the amount of data to be transferred and other factors such as network speed and disk I/O, this could take minutes or hours. Check the LUI logs to verify the successful completion of the transfer before proceeding to the next step. Logs are documented in Section 13 [Maestro Logs](#). Do not shut down the server while the transfer is taking place or you will have to restore from backup and start over.
8. If desired, delete the old LISTSERV Maestro database from the original database application. If the default internal database was previously being used, it can be disabled to save system resources. See Section 10.4 [Removing and Adding the Internal Database](#) for further instructions.

Under normal operation, it will never be necessary to revert from an external system database to the internal system database. However, if the need ever arises, follow these steps to restore the internal system database:

1. If the internal database was disabled, re-enable it. See Section 10.4 [Removing and Adding the Internal Database](#) for further instructions.
2. On the System Database Connection screen, select the **Use the internal database as the System Database** option.
3. Shut down LISTSERV Maestro and then restart.
4. Wait for the transfer of data to complete before deleting the external database.

10.2 Preparing the System Database

Before the Maestro User Interface can be used together with a freshly installed system database, the database must be prepared in certain ways. Outlined below are explanations of the required preparation steps, followed by details specific to each of the supported database management systems.

10.2.1 General System Database Preparation

LISTSERV Maestro must have its own database, separate from any other databases. The database may use the same database server as another database, but must not interact with the other databases on that server. Even if a database is created for storing recipient information for LISTSERV Maestro or for storing LISTSERV lists, it should be a separate database¹. LISTSERV Maestro can be given access to these recipient databases separately (see Section 9 [Defining External Database Connections](#)).

A user account must be created within the database server for LISTSERV Maestro to access the system database. This user will then be selected from the Maestro User Interface to connect to the database. Certain privileges are required for the user account, as described in the DBMS-specific sections below.

The database that is used as storage for the Maestro User Interface should be configured in a way that it allows dynamic growth because the data stored by the Maestro User Interface grows over time. The growth rate corresponds to the number and the size of the

1. This does not refer to the LISTSERV Maestro Data Warehouse, which stores and manages the hosted recipient data within the System Database. Instead, it refers to any recipient databases that were set up outside of LISTSERV Maestro (external user databases) from which LISTSERV Maestro may retrieve recipient data for email jobs.

email jobs that are delivered. Large email jobs with a high volume of collected tracking events will use more database storage space than smaller email jobs.

Some examples of upper limits that might need to be adjusted for large volume environments are:

User space quota – Most databases limit the amount of space that a given user may store in the database. This limit should be set to "unlimited" or a sufficiently large value for an organization's database usage.

Database or tablespace size – Many database vendors, especially those supporting larger database environments, support the sub-division of the database server in smaller areas, sometimes called "tablespaces" or a similar term (see the database documentation for details). Normally, each database account is assigned to one of these areas, which is then referred to as "default tablespace" or "standard tablespace". This part of the database should be configured in a way that it allows dynamic growth, if possible.



Note: It is possible to use the Maestro User Interface with a database that does not support this type of dynamic growth. To do so, an administrator should make it part of the daily or weekly routine to check the amount of space available for the Maestro User Interface, and then increase it manually when necessary.

File system size – Like other server applications storing persistent data on the file system, the database storing the Maestro User Interface data must reside on a server whose file system is monitored on a regular basis, either through automated system administration tools or by an administrator who regularly checks the system.



Important: Deleting or archiving old jobs from LISTSERV Maestro on a regular basis will prevent the database from becoming unnecessarily slow.



Tip: The amount of table space and file system space utilized by LISTSERV Maestro will vary drastically based upon the particular usage of the system. Regular monitoring of the database and disk usage are essential to ensure adequate storage space at all times.

10.2.2 Preparing SQL Server as the System Database

In the SQL Server Management Console, create a new database for sole use by the Maestro User Interface. Please see the SQL Server documentation for details about how to create, configure, and optimize a database.



Important: No matter if you create a new database or use an existing one, make sure that the database uses a *case-insensitive* collation. L-Soft recommends using the `Latin1_General_CI_AS` collation. This means that if you want to use an existing database, then you can only do so if this existing database already uses a case-insensitive collation.

Once a new database has been created, create a user account that the Maestro User Interface can use to connect to the database. Create a new user with any desired name and give it the `db_owner` role for the created or selected database.

10.2.3 Preparing Oracle as the System Database

A new Oracle database for sole use by the Maestro User Interface must be created so that it uses `UTF-8` as its database character set. The database character set `UTF-8` is

required and the Maestro User Interface will not work with a database that has a different character set. (See the Oracle documentation for details).

Use an Oracle administration tool (such as SQL*Plus), to create a new user. This new user must have the CREATE SESSION and the CREATE TABLE privilege and a sufficiently large table space quota in the user's default table space.

The Maestro User Interface does not require unusually large rollback segments. If duplicate elimination is performed for large email jobs, larger temporary segments are needed as duplicate elimination is performed with a database sorting operation. See the Oracle documentation for more details on how to configure and optimize databases.

The “maximum key length” value is a feature specific to Oracle. This value is an internal value inherent to each Oracle installation. It is determined mainly by the block size used by the database but may also be influenced by other factors, like the operating system.

For LISTSERV Maestro to be able to create its database table in an optimal manner, it needs to know the maximum key length value used by the Oracle database that is used as the system database. LISTSERV Maestro cannot query the database for this value. The administrator has to determine the maximum key length value used internally by the Oracle database installation and input the correct value.

If the value entered exceeds the actual maximum key length used by the database, runtime errors could result, and LISTSERV Maestro will not work correctly. If a value that is smaller than the actual value is entered, LISTSERV Maestro will tailor its database tables accordingly in order to meet this smaller value. As a result, the database tables will be created with a sub-optimal structure and the user may run into database column size limitations, which would be avoidable if the correct maximum key length value had been supplied on the System Database Connection Details screen.

Oracle documentation concerning which maximum key length value is used under which circumstances is sparse. A commonly used “rule-of-thumb” value is that with a block size of **4K**, the maximum key length is **1578**. With a block size of **8K**, the maximum key length is double the value of 4K, or **3156**. Generally, the maximum key length seems to be about 38% - 40% of the block size. If in doubt, please consult the Oracle documentation or contact Oracle support.

10.2.4 Preparing DB2 as the System Database

Use the DB2 Control Center application to create a new database and define UTF-8 as the database code set. The database code set UTF-8 is required and the Maestro User Interface will not work with a database that has a different code set. (See the DB2 documentation for details.)

Configure the database with at least one user table space and one system temporary table space with a page size of 32K each. It may also be necessary to create a buffer pool with a page size of 32K before the table spaces with 32K page size can be created. Next, create a new database user for sole use by the Maestro User Interface. This user must be configured to use a table space with 32K page size; otherwise the Maestro User Interface will not work. The new user needs the `Create Table` privilege.

If you are using DB2 V8.2, it is recommended that you use the V8.2 Thin Driver. If this is the case, then no further preparation is necessary.

If you are using DB2 V7.2, then the native driver is the only available driver. For the native driver, you need to create a database alias on the server that is running the Maestro User Interface component. In the Client Configuration Assistant, click the **[Add]**

button to create a new alias. Select the **Search the network** option and continue with the wizard. For more details on IBM DB2 database administration and the definition of database aliases, see the IBM DB2 documentation.



Notes: The name of this alias is the value for the **Database name** parameter of the IBM DB2 database plugin that comes with the Maestro User Interface.

The DB2 V8.2 database can also be accessed with the native driver; however, the V8.2 Thin Driver is highly recommended.

10.2.5 Preparing MySQL as the System Database

LISTSERV Maestro supports the older 3.23.x versions of MySQL (with 3.23.42 being the oldest officially supported build), as well as the current 4.0, 4.1 and 5.0 versions (5.0 being the most recent one at the time this was written). However, even though the older builds are supported, if you want to use MySQL as the system database, 4.1.7 or later is recommended.

To use the Maestro User Interface with MySQL, set up MySQL to use the InnoDB Tables table type (see the MySQL manual for more details). This table type supports transactions, and the Maestro User Interface requires a table type that supports transactions. Recent versions of MySQL (for example, version 4.1.7) come with a MySQL Server Instance Config wizard that lets the user specify settings for the database server. These settings will automatically be written to the MySQL configuration file. Within the configuration wizard, specify a database usage type that enables the transactional InnoDB storage engine. If no configuration wizard is available for your MySQL version, the InnoDB storage engine must be enabled manually in the configuration file. See the MySQL manual for more details.

In MySQL versions 3.23.50 and later, the InnoDB tablespace files can be configured to be “auto-extending”, i.e., they will grow automatically as needed. (This is the default when using the configuration wizard; otherwise, the auto-extend option should be manually specified in the configuration file.) For sites running earlier versions of MySQL, the InnoDB data files and log files must be created with sufficient size to accommodate the planned usage of LISTSERV Maestro. The database administrator should monitor the remaining capacity of the tablespace regularly and extend it as necessary.

After the MySQL database server is set up to use the InnoDB table type, create a new database specifically for use with the Maestro User Interface. Create a user to use with LISTSERV Maestro.

To connect to the database, start the MySQL client program `mysql.exe` from the `bin` folder of the MySQL binary installation (on Linux/Solaris, the client program is `mysql`.) To create a new database, enter the following command in the MySQL client: `create database DBNAME;` where `DBNAME` is replaced with the name of the database. Grant privileges by entering the following grant command for the username: `grant all on DBNAME.* to NAME@HOST identified by 'PASSWORD';` where the uppercase values are replaced as follows:

- **DBNAME:** The name of the database to be used with the Maestro User Interface. (This is usually the same name used in the “**create database**” command; see above. It is recommended that you use an all-lowercase spelling for the database name.)
- **NAME:** The user name of the user to be created and granted privileges.

- **HOST:** The host name of the server where the Maestro User Interface is running that will access this database.
- **PASSWORD:** The password associated with the user name.

Using “grant all” as described above grants all privileges on the given database to the given user. This is usually acceptable if the particular database was created specifically for use with the Maestro User Interface. However, if there is concern about granting the full set of privileges to the user, use the following privilege list instead of “all”:

```
select, insert, update, delete, index, create, drop
```

10.3 General Optimization Hints for the System Database

The following general information about how the Maestro User Interface uses the database can help optimize the database installation for use with the Maestro User Interface.

- The Maestro User Interface does not use large transactions. Any transactions that are opened are then closed after a maximum of a few hundred inserts or updates.
- During normal usage, the Maestro User Interface behaves with OLTP (online transaction processing) characteristics. There is a constant switch between read and write on the database. However, if there are many reports running on the collected tracking data, the characteristics of the Maestro User Interface’s behavior shift more and more into OLAP (online analytical processing), where the amount of (complex) reads outnumbers the amount of writes.

Use this information to optimize the database after analyzing the usage of the Maestro User Interface to determine if it is working more with OLTP or OLAP characteristics.

10.4 Removing and Adding the Internal Database

If you connect the Maestro User Interface to an external database, there is no longer a need for the internal MySQL database to run together with LISTSERV Maestro. To reduce the resource usage of LISTSERV Maestro, disable the internal database after configuring the Maestro User Interface to use an external database.

To uninstall the internal database, you’ll need to run the LISTSERV Maestro Setup wizard on the server where the Maestro User Interface component is installed. (See the Installation Manual for further information.) Once you start the Setup wizard, simply **modify** the installation and deselect the **Internal MySQL database** component. Leave the other components as they were. The Setup wizard will automatically uninstall the internal database.



Important: After the internal database is disabled, it is no longer possible to select the **Use the internal database as the System Database** option on the System Database Connection screen of the Maestro User Interface. If this choice is selected by mistake, and LISTSERV Maestro is restarted with no external database configured, the Maestro User Interface component will not start, as it will not be able to find the now disabled instance of the internal database. If, at a later point, the internal database is desired, it must be re-installed.

Section 11 Saving and Restoring a Backup

It is a standard best practice for any administrator to make regular backups of critical software and data. LISTSERV Maestro archives a consistent backup of the data collected in the application so that it can be restored in the event of a system failure.

LISTSERV Maestro gives the Administration Hub component the responsibility of acting as backup master to avoid any problems that might arise from having different components that store data independently and reside on different servers. If different components initiate backups at different times, inconsistent data sets between components can result. If both backups were then to be restored, the data sets would be inconsistent, invalidating the backup.

The Administration Hub will centrally trigger a backup on all connected components (including itself) in order that the backup data saved by each component is consistent with the backup data of all other components. This backup is initiated based on the values entered in the Global Component settings for the Administration Hub.

If regular backups are performed through LISTSERV Maestro, no external backups of the Maestro database are necessary. In the event of a system crash, the data stored in the Maestro database would not be sufficient to restore the Maestro application, as the data would not be synchronized with Maestro's internal registry. Restoring LISTSERV Maestro can only be done from a complete backup initiated from within LISTSERV Maestro. Performing an external backup of the Maestro database serves no useful purpose and may even cause problems when Maestro attempts to access a database that is in backup mode.

The backup procedures described here are specifically for LISTSERV Maestro. If there are other databases on the same server as the LISTSERV Maestro database, they should be backed up separately.



Note: The procedures described here refer only to the three Maestro components – HUB, LUI, and TRK. If other elements are located within the LISTSERV Maestro tree (i.e. the LISTSERV web interface files or independent Web pages served up by Tomcat), then these must be backed up and restored independently.

11.1 Configuring the Backup Time

The application wide backup is triggered once per day. Each day at a certain time, the Administration Hub (backup master) will start a backup of each component. To assign backup settings, click on the **Global Settings** menu, select **Administration Hub**, and then finally select **General Administration**. In the **Time for daily backup** field, set the time to start the backup master by entering the desired time of the daily backup in the form of hh:mm with values from 00:00 to 23:59.

There may be times when it is necessary to create a backup immediately; for example, just before any invasive procedure such as moving a component or applying a patch, or in the case of an emergency. To perform a backup immediately, click the **[Execute Backup Now]** button. Click **[OK]** to save settings and return to the Administer Component Settings screen.



Note: The **[Create Test-Bed Backup]** button should not be used for regular backups. Its use is documented in Section 12 [Using a Test-Bed Backup](#).

Figure 11-1 The General Component Settings for Administration Hub Screen

General Component Settings for Administration Hub

Backup Settings

Time for daily backup: (hh:mm)
(This time is relative to the time zone of the server)

External processes to run after a backup completes:
(Your commands will be executed inside of a command shell with "cmd /c YOUR_COMMAND ARG1 ... ARGn".)

After backup success: [Test](#)
(Leave empty for default: "none".)

After backup failure: [Test](#)
(Leave empty for default: "none".)

Work folder for commands:
(Leave empty for default: "application home".)

[View list of currently active "after backup" processes](#)

Backup folder:
(Leave empty for default: "backup".)

Number of previous backups to keep:

Advanced Security Options

Limiting the number of allowed unsuccessfull login attempts is a measure to avoid external automated dictionary attacks. Supply the maximum number of allowed unsuccessfull login attempts below to activate this advanced protection.

Maximum Number of Unsuccessfull Logins: (leave empty for unlimited number)

Lock Login Time in Minutes:

11.2 Configuring External Post-Backup Processes

The administrator may define external processes that will be executed after a backup is completed. External processes may be used to execute additional backup tasks such as automatically moving the backup folders to a tape, copying backup folders to a network drive, notifying the administrator by email if the backup was unsuccessful, and so on. Two different external processes can be defined, one to be executed after a successful backup and one after a backup failure.

Each process is specified in form of an external command that is executed by the Administration Hub when the backup completes. If it is necessary to execute more than one command, they can be written into a batch file (Windows) or shell script file (Unix/Linux). If this is the case, the name of that batch/script file is entered as the external command to be executed (with all necessary parameters). The administrator may also specify the work folder for the commands (same folder for both commands).

Clicking **Test**, located next to each command box, executes the command for testing. A new window will pop up that shows the output of the command. In this window, the external process can be stopped, if necessary. Closing the window before the process terminates, will not stop the process. To view the output of the test process again (if it is still running), or to terminate it (if it does not terminate by itself), access the process by using the **View list of currently active "after backup" processes** link.



Note: Commands must *not* define external processes that run indefinitely. Each external process should terminate itself when it has completed the action. Processes that run continuously slow the server down and will cause a crash because each time a backup finishes, a new process will be started, tying up more system resources.

If several external processes are running using a batch/script file, make sure that all processes started by the batch/script file terminate themselves at some point. If an external process that does not terminate itself is started, (because of a defect in the external process, or by mistake) click on the **View list of currently active "after backup" processes** link.

This screen displays a list of all currently active external processes started by the Administration Hub, either as an actual "after backup" process that was started when a backup was completed, or because the administrator clicked on one of the **Test** links. Only processes that are still running are shown in this list. Each process is shown with the date and time it was started, the command that was used to start it, and a link that opens a pop-up window. The pop-up window continuously shows the output of the external process (if any) and allows for the termination of that process while it is still running.

If any of the command fields are left empty, no external process will be started at the corresponding "after backup" condition. If the work folder is left empty, then the application home folder of the Administration Hub will be used as the work folder.

11.3 Configuring the Backup Location

Each component has a backup location. This is necessary because the components may reside on different servers. The backup default location is the `backup` folder, which is in the `home` folder of the component in question (for example "`\Program Files\L-Soft\Application Server\lui`").

It is possible to use a different folder if desired. The folder configured may be either an absolute path, such as "`C:\MyFolder\backup`", or a relative path, such as "`myFolder\backup`", which is then interpreted as being relative to the home folder of the component. Enter the path name in the **Backup folder** field at each of the following locations:

- For the Administration Hub component, click on the **Global Settings** menu, select **Administration Hub**, and then finally select **General Administration**. The General Component Settings for Administration Hub screen opens.
- For the Maestro User Interface component, click on the **Global Settings** menu, select **Maestro User Interface**, and then finally select **General Administration**. The General Administration of Maestro User Interface screen opens.
- For the Maestro Tracker component, click on the **Global Settings** menu, select **Maestro Tracker**, and then finally select **General Administration**. The General Component Setting for Maestro Tracker screen open.



Important: Do not configure different components to save backups into the same folder. Doing so may cause backups from one component to over-write backups from another, resulting in data loss. Each component must have its own dedicated backup folder.

11.4 Configuring the Backup History

To lessen the risk of restoring a backup containing corrupted data, LISTSERV Maestro provides the opportunity for administrators to create a backup history. Each time a new backup is made, it is saved into the backup folder configured for the component (see Section 11.3 [Configuring the Backup Location](#)).

If the component is also configured to keep a number of previous backups, then the folders containing the older backups will be kept under names like “NAME1”, “NAME2” ... “NAME_n”, where “NAME” is the name of the standard backup folder and “n” is the number of previous backups that the component is set to keep.

For example, if a component is configured to keep three previous backups, then the backup history of each day will look like this:

Table 11-1 Backup History

Day 1	backup – contains backup of day 1
Day 2	backup – contains backup of day 2 backup1 – contains backup of day 1
Day 3	backup – contains backup of day 3 backup1 – contains backup of day 2 backup2 – contains backup of day 1
Day 4	backup – contains backup of day 4 backup1 – contains backup of day 3 backup2 – contains backup of day 2 backup3 – contains backup of day 1
Day 5	backup – contains backup of day 5 backup1 – contains backup of day 4 backup2 – contains backup of day 3 backup3 – contains backup of day 2

Keeping a backup history can help ensure against corrupted backup data. However, as the amount of application data grows, it may not be possible to keep many old backups, which take up space on the disk. In addition, keeping older backups on the same disk does not ensure against failure of the disk itself (head crash for example). *Always* save the backup to an external backup medium as described in Section 11.5 [Saving a Backup to an External Medium](#).



Note: If daily backups are saved to an external medium routinely, it is acceptable to set the number of old backups to “0”.

11.5 Saving a Backup to an External Medium

Once LISTSERV Maestro has completed its backup, the configured backup folder of each component contains the data that is required to restore this component to the state

of the moment when the backup was triggered. To prevent catastrophic loss of data, save these folders to an external backup medium (e.g. a backup tape or storage device).

To avoid a potential partial backup problem, either use the automatically triggered external post-backup process, outlined in Section 11.2 [Configuring External Post-Backup Processes](#), to ensure that the backup tool does not start its work until after the completion of the internal backup (recommended), or use whatever standard backup tool is used by the organization to configure a daily backup of the designated folders. Schedule this daily backup to occur after the time when the Administration Hub itself completes the backup of each component. There should be a enough time between the backup triggered inside of LISTSERV Maestro and the backup to the external medium triggered by the backup tool to ensure that all components have enough time to complete their backups. Otherwise, partial data would be backed up to the external medium.

For small installations, the backup inside LISTSERV Maestro will not take more than a few minutes. However, as the data in the LISTSERV Maestro installation accumulates over time, backup naturally will take longer. If post-backup triggers are not being used, periodically check the backup logs to see how long the backup actually takes and schedule the external backup accordingly, at a safe time after the LISTSERV Maestro backup is completed.

Remember that the external post-backup command or backup tool must be configured such that it backs up all backup folders of all components. A LISTSERV Maestro installation will have three backups to save to an external medium, one for the Administration Hub, one for the Maestro User Interface, and one for Maestro Tracker. These folders may also reside on different servers, depending on the installation.

11.6 Identifying the Backup: The Backup ID

Because the LISTSERV Maestro components store their backup data into separate folders, it is necessary to know which of the folders belong together, in case a backup history is kept or it is necessary to retrieve a backup from an external medium. This is done using the backup ID. Each backup gets a unique ID that is shared by all components participating in the backup. Each component also writes a “`readme.txt`” file into the backup folder. Stored in this text file is the ID of the backup that saved the data in the particular backup folder, together with output about backup start time, end time and its success or error state.

11.7 Restoring a Backup

In the unfortunate event of having to restore a backup, follow the procedure described in this section. Several steps need to be executed to restore a backup successfully. Please review each step carefully. Some steps have lengthy descriptions or sub-steps. Do not skip steps or do them out of order, or the restoration will not succeed.

1. Identify the backup that is to be restored.

This usually is the most recent backup, but it also must be a successful backup. If there were errors during the most recent backup, revert to the next most recent backup, which may have to be retrieved from an external medium. To find out if a backup was successful check the backup log in the folder:

```
[maestro_install_folder]/hub/logs
```

For each backup triggered by the Administration Hub, a report named “backupReport_ID.txt”, where “ID” is replaced with the ID of the backup in question, is saved into this folder. The IDs are assigned in alphanumeric order; the most recent backups have higher order IDs (in an alphanumeric sense). Use the file date of the report file to locate the most recent backup. If the backup was successful, an entry like this will appear at the end of the file:

*“The backup was completed successfully
Final completion date: <date here>”*

If the backup was unsuccessful for any reason, then the report will contain entries detailing the errors that occurred.

If the logs folder cannot be accessed, (because a disk crash destroyed the disk of the Administration Hub installation, for example) it is still possible to locate the most recent successful backup by opening the `readme.txt` file in the backup folder of each component. The `readme.txt` file lists the backup ID and the success state of that particular backup. If no errors are reported in that file, then the backup of this component was successful. If successful backups with the same ID of all the other components are located, then a complete and successful backup set exists and can be restored.

2. Find the backup folders from all components that belong to the same backup set.

Once the backup to be restored has been identified and the backup ID is determined, the next step is to find all the backup folders of the individual components that contain data for this backup. Check the `readme.txt` in the backup folder of each component. If it contains the same ID, the right backup folder for this component has been located.

3. If necessary, make a fresh installation of all components.

If the original system is still in working order and the purpose of the backup restoration is simply to restore a previous state of the application user data (for example if the application user data was corrupted, or if some LISTSERV Maestro objects, like a mail job or hosted recipient data, were accidentally deleted), then it is not necessary to do a reinstallation of the application. In this case, simply proceed with the next step.

If the original system was destroyed (for example by a disk crash) or generally damaged (where it is not clear if the damage is limited to the application user data or may also have affected the application binaries), then you will need to do a fresh installation (this also includes uninstalling the old files, unless you start from scratch on a new system). If you do this, install the components on the servers where you need them. After the re-install, do not start the components. Instead proceed as described in the next step.



Note: If LISTSERV Web Interface files and/or other non-Maestro files are maintained within the Maestro application folder tree, then care should be taken to preserve them, if possible, before wiping out the old installation.

4. Restore all three components.

To restore the Administration Hub:

Remove the existing versions of the file “hub.ini” and the folders “accountreg” and “hubreg”, including their contents, from the Administration Hub home folder:

```
[maestro_install_folder]/hub
```

Replace them with the versions from the backup folder of the Administration Hub component.

To restore the Maestro User Interface:

Remove the existing versions of the files “lui.ini” and the folders “luidata” and “registry”, including their contents, from the Maestro User Interface home folder:

```
[maestro_install_folder]/lui
```

Replace them with the versions from the backup folder of the Maestro User Interface component.



Important: If the backup is from a LISTSERV Maestro version earlier than 2.1, the backup may also contain a file called “my.ini”. This file is no longer required by LISTSERV Maestro 2.1 and should not be restored from the backup.

Next, use a text editor (i.e. Notepad on Windows) to add a new entry into the “lui.ini” file like the example below:

```
RestoreBackup=path_to_backup_folder
```

The “*path_to_backup_folder*” is replaced with the path name that leads to the backup folder from which the files and folders, as described above, were copied.

This path name may either be a full path name including drive letter, or it may be an absolute path without drive letter starting with “\” or “/”, which is then interpreted as being absolute on the drive/root where the application server is installed (for example, in the default case for Windows, the same drive where “\Program Files\L-Soft\Application Server” is located). Or a relative path without a driver letter may be used, and not starting with either “\” or “/”, which is then interpreted as being relative to the home folder of the Maestro User Interface component (for example, in the default case for Windows, that would be the folder “\Program Files\L-Soft\Application Server\lui”).

Forward slashes “/” or backslashes “\” may be used as the filename separator. However, if backslashes are used, then use double backslashes.

Example, either write:

```
C:/Sample/MyFolder/backup
```

– or –

```
C:\\Sample\\MyFolder\\backup
```

This entry to the “`lui.ini`” file will be automatically removed during the first startup of the component. It is only present to signal to the component that it should restore all required data from the given folder, which happens automatically during the next startup, whenever this INI file entry is present. For more information on editing INI files, see Section 20 [Editing LISTSERV Maestro INI Files](#).

To restore Maestro Tracker:

Remove the entire `data` folder from the Maestro Tracker home folder:

```
[maestro_install_folder]/trk/data
```

Replace it with the version from the backup folder of the Maestro Tracker component. Also remove the `tracker.ini` file in the Maestro Tracker home folder, and then replace it with the same file from the backup folder.



Note: If the backup is from a LISTSERV Maestro version earlier than 2.0-4, then the backup may contain several `*.dat` files instead of a single `data` subfolder, which was introduced in 2.0-4. In this case, restore the backup as follows:

Remove all `*.dat` files from the `data` folder inside of the Maestro Tracker home folder:

```
[maestro_install_folder]/trk/data
```

Replace them with the `*.dat` files from the backup folder of the Maestro Tracker component.

5. Edit respective INI files, if necessary.

If components are being restored on different servers or a different combination of servers than where the original backup was taken from, it may be necessary to edit the respective “`*.ini`” files of the components. This would include restoring a backup to a server with a different name, using a different port number, or changing how the components are grouped on a server or servers. For example, if components that were all originally on the same server are moving to different servers, or taking components that were originally on different servers and moving them to the same server.

If LISTSERV Maestro is using a default internal MySQL database that has undergone modifications or optimizations to its configuration (because of changes made through the MySQL configuration tools or by manual edits to the “`my.ini`” file in the “`lui\database`” folder), those modifications must be re-applied. The freshly installed LISTSERV Maestro contains an internal database with the default configuration.

6. Check these settings.

Check the settings in

```
Windows: [maestro_install_folder]\commands\UserConfig.cmd
```

```
Unix/Linux: [maestro_install_folder]/commands/UserConfig.sh
```

to verify that all settings (like minimum/maximum memory allocation, run-user, etc.) are correct for the installation. If you still have access to the corresponding file of the original installation, you can compare with the values in that file.

7. Restore other files, if necessary.

If LISTSERV Web Interface files or any other non-Maestro files were stored in the Maestro folder tree, then restore them to their proper location.

8. Start all components.

During startup, the system database content will be restored from the backup folder. Monitor the log files of the components to check if they start up correctly. If yes, the backup restoration is complete. If any component does not start up correctly, this may be because of differences in the configuration of the backed up system and the restored system. In that case, it may be necessary to adjust further INI file settings (see previous step) or to log into the Administration Hub and configure the necessary settings accordingly. Then restart and again monitor the startup log entries. If necessary, repeat this until the system starts up normally.



Note: If you are using the internal MySQL system database and have previously made modifications or optimizations to its configuration (for example by using the MySQL configuration tools or by manually editing the file “my.ini” in the folder “lui/database”), then you now need to re-apply the same modifications, since the freshly installed LISTSERV Maestro once again contains an internal database in the default configuration.

Section 12 Using a Test-Bed Backup

In certain situations it may be useful to make a copy of all the data in a given LISTSERV Maestro installation and transfer it into a second (test-bed) installation without affecting the production installation.

Here are a few situations where such a test-bed installation might be needed:

- **Testing a software upgrade** – If the production LISTSERV Maestro installation has a high-availability requirement, then it is prudent to upgrade a test server first to make sure the upgrade will not result in unanticipated down-time of your production system. L-Soft performs thorough testing before releasing new versions, but it is not always feasible to test every possible situation. To make sure the upgrade process will work flawlessly with your data, you can set up a “test-bed” server with your production data and perform the upgrade there first. If any errors result, these can be addressed by L-Soft support before you upgrade your production server.
- **Beta testing** – If you are participating in a beta-test of an upcoming version of LISTSERV Maestro, then you may want to test the new features using your production data, but without running beta software in production. By beta-testing in an environment identical to your production environment and using your production data, you can ensure that there won’t be any issues specific to your installation when the product is released. By trying out the new features with your production data, you can anticipate how you might use those features in production and even suggest improvements to the developers so that the new features really meet your needs. It is much easier to discover such opportunities for improvement when working with real data and realistic scenarios than when using test data and test scenarios.
- **Training** – If you are conducting training (or hiring L-Soft to conduct training at your site) in LISTSERV Maestro for your users, then it may be helpful to conduct the training using real data and realistic scenarios.

In order to create a Test-bed backup, you cannot simply trigger a regular backup on the original system, and then restore this backup into the test system. Doing so has several pitfalls which could, in a worst case scenario, destroy some of the data in the original production system.

Some of the possible pitfalls of using a regular backup to create a test-bed are:

- If, at the moment the backup is triggered, the original system contains a mail job in the Outbox and it’s scheduled for delivery, and this backup is restored into the second system, and the scheduled delivery time arrives, then the second system would actually send out the job to the given recipients.
- And since the job is a copy of a real job on the production server, with real recipients, this mailing from the test system would go out to these real recipients. At the same time, since the same job also still exists on the production server, it will be

mailed out from there. As a result, the recipients will actually get two copies of the same mailing.

- If the production server installation is distributed over several servers (for example where LUI is installed on one server and TRK is installed on another server), then the information about how the components are distributed will also be contained in the backup data. Then, you would have to be very careful when restoring this backup into the test server so that you do not accidentally end up, for example, connecting the test server LUI with the TRK of the original production system (or vice versa).
- If the production server installation uses an external system database, then the connection information for this external database is also stored in the backup data. This means that, during the backup restoration, this information will also be restored to the test server. Because of this, the test server will try to connect to the same external database (and use it as its system database) as the production server. At worst, the test server could then delete or change data belonging to the production system!

Because of these pitfalls, it is generally not advisable to restore a backup of an existing and running system into a different system. Therefore, the idea of restoring a backup from the original server to the test server is not a good idea.

To solve this problem, LISTSERV Maestro offers the *Test-Bed Backup feature*. A test-bed backup is similar to a normal backup of a system, meaning it contains all the system data (all user accounts, jobs, reports, tracking information, hosted recipient data, etc.). However, all the critical aspects of the data, like information about distributed components, connection information for external databases, scheduled jobs in the outbox, etc. have been changed by the system (when the data was written to the test-bed backup) so that they no longer pose a risk when the test-bed backup is restored to a test server. Simply put, a test-bed backup is as close a copy to the data of the original system as possible, but with all the critical information removed or changed.

Information that is removed or changed includes:

- **Component distribution information.** If the original system was distributed over several servers, then the test-bed backup is no longer aware of this and assumes that all components are on the same server (i.e. the “localhost”).
- **External database information.** If the original system used an external system databases, then the test-bed backup is no longer aware of this and assumes that the internal system database is used.
- **Instance-ID information.** The test-backup backup does not contain the original system’s instance ID. This means that during the first start of the test server, it will generate its own instance IDs.
- **LISTSERV connection host information.** The test-bed backup does not contain information about which LISTSERV host to use; therefore, this information needs to be added manually (via the Administration Hub) after the test-bed backup has been restored to a test server. This allows the test server to connect to a different instance, if desired.

- **Tracking host information.** The test-bed backup does not contain information about which tracking host name to use; therefore, this information needs to be added manually (via the Administration Hub) after the test-bed backup has been restored to a test server.
- **Scheduled jobs.** The **outbox send queue** option in the test-bed backup is set to **Sending is disabled**. This means that any queued jobs in the Outbox are not automatically delivered by the server into which the test-bed backup is restored. If you need to send any jobs on the test server, then it is therefore necessary to re-enable the Outbox (via the Administration Hub) on the test server. However, before re-enabling the Outbox, make sure to check it for any production jobs still left in it. If there are any left, then delete or revoke those jobs to stop them from being delivered when the Outbox is re-enabled.

With the risky information being removed or reset to appropriate defaults, it is safe to restore the test-bed backup to a test server without the risk of damaging or impacting the production server.

The following sections describe how a test-bed backup can be created on an original system and how it must be restored into a test system.

12.1 Creating a Test-Bed Backup on the Original System

To create a test-bed backup on the original system, log into the Administration Hub of the original installation. Click on the **Global Settings** menu, select **Administration Hub**, and then finally select **General Administration**. From the General Components Settings for Administration Hub screen, click the **[Create Test-Bed Backup]** button.

Once triggered, the test-bed backup proceeds just like a normal backup, with just a few differences:

- The folders where the backup data is stored are always the following three folders:
 - For the HUB component: `[install_folder]/hub/test-bed_backup`
 - For the LUI component: `[install_folder]/lui/test-bed_backup`
 - For the TRK component: `[install_folder]/trk/test-bed_backup`

The backup is always written to a folder called “test-bed_backup” inside of the home folder of each of the three components (on their respective server, if the components are installed on different servers).

- Just like with a normal backup, a test-bed backup consists of the backups of all three components, i.e. a complete test-bed backup encompasses all three of the above folders.

For each test-bed backup that is created, a backup report is written to the log folder of the Administration Hub component (“`[install_folder]/hub/logs`”), just like for a normal backup. The content of this report is also very similar to a normal backup report, only that it clearly states that this backup is a test-bed backup (and thus contains data which has been changed slightly, in comparison to the original data).

- Similarly to a normal backup, a test-bed backup is only valid (i.e. contains complete and restorable data) if the last line of the report reads:

The test-bed-backup was completed successfully

If you do not see this line, then the backup was not successful and should not be used during the restore step described in the next section.

12.2 Restoring a Test-Bed Backup into the Test System

A test-bed backup is restored exactly like a normal backup; see Section 11 [Saving and Restoring a Backup](#) for more information.

Just like a normal backup, it is important to restore *all three* components and also to make sure that the backup data for each component was written by the same backup (or test-bed backup in this case). As with a regular backup, this verification is done by verifying the backup ID (Test-Bed-Backup-ID) in the three backup parts.

The test-bed backup is not restored on the original system where it was created; instead, it is restored on a different system or the “test server” system. It is also important to know that a test-bed backup must only be restored to a test server that was installed on a single server, including the internal database. A test server must meet the following requirements:

- All three components (HUB, LUI and TRK) are installed on the same server.
- The internal MySQL database is also installed and used as the system database.

These requirements are easily met by installing a fresh installation of LISTSERV Maestro on the test server, and then selecting all three components, plus the internal system database, when running the Setup wizard (or using the **Express Setup** option in the Suite Installation Kit for Windows).

Once you have such a test server system and have also identified the test-bed backup you want to restore, then you can restore it by following the normal backup restoration steps.

Section 13 Maestro Logs

LISTSERV Maestro log files are located in two places. Log files having to do with specific LISTSERV Maestro components are kept in a directory configured like the example below:

```
x:\Program Files\L-Soft\Application Server\XXX\log
```

“x:” is the drive where LISTSERV Maestro is installed and “xxx” is the component, either HUB, LUI, or TRK.

Log files for third party components like Tomcat are kept in a directory configured like the example below:

```
x:\Program Files\L-Soft\Application Server\logs
```

“x:” is the drive where Maestro is installed.

13.1 Remote Log Access

The three main LISTSERV Maestro components all write their own log files. These files can be found in the `logs` subfolder of each component’s home folder inside of the `Application Server` installation folder. In some situations, the Maestro administrator may not have access to these folders, but still needs to access the log files. To solve this, LISTSERV Maestro offers remote log file access. The remote access allows the Maestro administrator to download the log files from the server through a web browser.

Before accessing the log files of a component, configure the component for remote log access. To do so, edit the INI file of the component and add the following entry:

```
RemoteAdminPassword=PASSWORD
```

Replace `PASSWORD` with a password known only to authorized administrators. For security reasons, do *not* use the normal admin password from the Administration Hub. Because this password will later be used as a parameter in a URL, use only URL-safe characters in the password (alphanumeric characters).

Remember; add this entry to *each* component’s INI file; to `lui.ini`, `hub.ini` and `tracker.ini`. For information on how to edit INI files, see Section 20 [Editing LISTSERV Maestro INI Files](#). If the entry is not added to one of the INI files, then it will not be possible to access the log files of that component, but it will still be possible to access logs of the other components where the entry has been added. To disable remote log access, simply remove the entry from the INI file(s) or comment it out. Whenever this entry is changed, the change will be effective *immediately* – The component will *not* have to restart.

Once the component(s) have been configured for remote log access, access their log files from any web browser on any computer that has HTTP access to the particular component. The only requirements for access are the `PASSWORD` configured in the INI file(s) and the date of the log file to access.

- To download a **Maestro User Interface** log file, access the following URL:
`http://HOST:PORT/lui/downloadLog?pw=PASSWORD&day=DATE`

- To download an **Administration Hub** log file, access the following URL:
`http://HOST:PORT/hub/downloadLog?pw=PASSWORD&day=DATE`
- To download a **Maestro Tracker** log file, access the following URL:
`http://HOST:PORT/trk/downloadLog?pw=PASSWORD&day=DATE`

Replace `HOST` with the host name of the server running the component to be accessed, `PORT` with the HTTP port on that server (`:PORT` can be left out if the HTTP-port is 80), `PASSWORD` with the password configured in the INI file, and `DATE` with the date of the day of the log file to download. The date is formatted as `YYYYMMDD`, where `YYYY` is the year with 4 digits, `MM` is the month with 2 digits and `DD` is the day of the month with 2 digits.

13.2 Subscriber Activity Change Log

LISTSERV Maestro offers the option of keeping a change log of all subscriber activities (i.e. a log of all subscribe, unsubscribe, join, un-register, and address-change activities of subscribers and members of hosted lists and datasets). By default, the subscriber change log is deactivated.

The change log is activated with the following entry in the `lui.ini`:

```
ChangeLog=true
```

Optionally, a change log time period can be specified to determine how often a new change log is started. Possible time periods are `daily`, `weekly`, `monthly`, and `yearly`. The time period is specified after the keyword `true` in the INI-file, separated with a comma, similar to this example:

```
ChangeLog=true,daily
```

If no time period is specified, the default `weekly` will be used. This means that:

```
ChangeLog=true
```

is equivalent to

```
ChangeLog=true,weekly
```



Notes: The keywords are not case sensitive, meaning that you could also type `True`, `Yearly` or `TRUE`, `MONTHLY`.

If any other value than `true` is specified, then the change log is deactivated. Therefore, you can temporarily deactivate the change log either by commenting out the `ChangeLog` line, or by changing the value `true` to something else, for example:

```
ChangeLog=false,daily
```

Each line in the change log corresponds to one subscriber or member activity. Each line is prefixed with the date and time of the activity, followed by a three letter activity code and activity details. The following activities are logged:

- `NEW D dataset_id email_address ip_address`

Logged when a new dataset member was added to a dataset:

- *dataset_id* – The ID of the dataset (an integer number).
 - *email_address* – The email address of the new member.
 - *ip_address* – Appears only if the new member was added because of an active join by an actual user; in which case, the user’s IP-address is logged. If the new member was added by the LISTSERV Maestro data admin, then no IP-address is logged.
- DEL D *dataset_id email_address ip_address*

Logged when a dataset member was deleted from a dataset:

- *dataset_id* – The ID of the dataset (an integer number).
 - *email_address* – The email address of the member that was deleted.
 - *ip_address* – Appears only if the member was deleted because of an active un-register by the actual user; in which case, the user’s IP-address is logged. If the member was deleted by the LISTSERV Maestro data admin, then no IP-address is logged.
- ADR D *dataset_id old_email_address new_email_address ip_address*

Logged when a dataset member’s email address was changed:

- *dataset_id* – The ID of the dataset (an integer number).
- *old_email_address* – The old email address of the member.
- *new_email_address* – The new email address of the member.
- *ip_address* – Appears only if the member’s address was changed because of an active change by the actual user; in which case, the user’s IP-address is logged. If the member’s address was changed by the LISTSERV Maestro data admin, then no IP-address is logged.



- TP_GRANTED *dataset_id email_address ip_address*

Logged when a dataset member grants permission for personal tracking (indicated by setting a dataset profile field with the special type “Tracking Permission” to “true”):

- *dataset_id* – The ID of the dataset (an integer number).
- *email_address* – The email address of the member.
- *ip_address* – Appears only if the member’s address was changed because of an active change by the actual user, in which case the user’s IP-address is logged. If the member’s address was changed by the LISTSERV Maestro data admin, no IP-address is logged.



- TP_REVOKED *dataset_id email_address ip_address*

Logged when a dataset member revokes permission for personal tracking (indicated by setting a dataset profile field with the special type “Tracking Permission” to “false”):

- *dataset_id* – The ID of the dataset (an integer number).
- *email_address* – The email address of the member.
- *ip_address* – Appears only if the member’s address was changed because of an active change by the actual user, in which case the user’s IP-address is logged. If the member’s address was changed by the LISTSERV Maestro data admin, no IP-address is logged.

- NEW L *dataset_id list_id email_address ip_address*

Logged when a new subscriber was added to a list:

- *dataset_id* – The ID of the dataset the list belongs to (an integer number).
- *list_id* – The ID of the list (an integer number).
- *email_address* – The email address of the new subscriber.
- *ip_address* – Appears only if the new subscriber was added because of an active subscribe by an actual user; in which case, the user’s IP-address is logged. If the new subscriber was added by the LISTSERV Maestro data admin, then no IP-address is logged.

- DEL L *dataset_id list_id email_address ip_address*

Logged when a subscriber was deleted from a list:

- *dataset_id* – The ID of the dataset the list belongs to (an integer number).
- *list_id* – The ID of the list (an integer number).
- *email_address* – The email address of the subscriber that was deleted.
- *ip_address* – Appears only if the subscriber was deleted because of an active unsubscribe by the actual user, in which case the user’s IP-address is logged. If the member was deleted by the LISTSERV Maestro data admin, then no IP-address is logged.



Notes: If the list is a hosted LISTSERV list and the user unsubscribed from the list by sending an “unsubscribe” email to LISTSERV, then the address “0.0.0.0” is logged, since the real IP-address of the user is not known in this context.

In addition to the above entries, the system also writes marker entries to the change log whenever a backup is performed (i.e. automatically once a day) or restored (i.e. when a backup-restore is initiated by the administrator).

Section 14 Using Non-Standard Ports

The components of LISTSERV Maestro use a number of ports to communicate with each other and with the external world. The ports used are standard ports and will work well under most circumstances. Under certain conditions, it may be desirable to change one or several of the ports to other ports – for example, if another application installed on the same server already uses one of the ports LISTSERV Maestro is set to use. Changing ports may require editing certain INI files. For more information on editing LISTSERV Maestro INI files, see Section 20 [Editing LISTSERV Maestro INI Files](#).

14.1 Ports Used by LISTSERV Maestro

This list contains the individual ports used (by default) by each of the LISTSERV Maestro components.

14.1.1 Ports used by the Administration Hub

The Administration Hub uses three different ports:

- For HTTP access to the Administration Hub user interface (using a web browser), the Administration Hub uses the standard HTTP port **80** (or **443** for HTTPS).
- For internal communication with the other components, the Administration Hub uses port **1099**.
- For shutdown of the application server, the Administration Hub uses port **8007**.

14.1.2 Ports used by the Maestro User Interface

The Maestro User Interface uses four different ports:

- For HTTP access to the Maestro User Interface (using a web browser), the Maestro User Interface uses the standard HTTP port **80** (or **443** for HTTPS).
- For internal communication with the Administration Hub, the Maestro User Interface uses port **1099**.
- For the internal database connection, the Maestro User Interface uses port **3306**.
- For shutdown of the application server, the Maestro User Interface uses port **8007**.

14.1.3 Ports used by Maestro Tracker

Maestro Tracker uses four different ports:

- To collect the tracking events from mailings sent with the Maestro User Interface, Maestro Tracker uses the standard HTTP port **80** (you cannot use HTTPS for Tracker).
- For internal communication with the Administration Hub, the Maestro User Interface uses port **1099**.

- To transfer the tracking events to the Maestro User Interface, Maestro Tracker, uses port **7000**.
- For shutdown of the application server, Maestro Tracker uses port **8007**.

14.2 Configuring Port Usage

If any of the ports described in the previous sections are already in use on the server where the LISTSERV Maestro component is installed, it is possible to change the use of this port. Note that some components make use of the same port as other components. This is not a problem between the different components of LISTSERV Maestro. If there are several components on the same server, then the components share usage of these ports (port 80 for HTTP access and port 1099 for internal communication, for example). It is not necessary or even possible to configure one component to use a different port than the other while the components are on the same server.

14.2.1 Configuring the HTTP Port

To configure the HTTP port, edit the `Port` entry in the Tomcat INI file:

```
[maestro_install_folder]/conf/tomcat.ini
```

Example: `Port=8080`

If you leave out the entry from the `tomcat.ini` file, the default port “80” will be used for HTTP connections.

If there are several LISTSERV Maestro components installed on the same server, then they will all be affected by this change. It is *not* possible to use different HTTP ports for each of the components if the components are installed on the same server. However, if the components are installed on different servers, they can use different HTTP ports. These changes will only be effective after a restart of the component in question.

When changing the HTTP port, there are a few issues of which to be aware:



Note: These changes will only become active after LISTSERV Maestro has been restarted, i.e. after you have performed all the steps described below, restart LISTSERV Maestro. (Except in those cases where it is explicitly noted that you need to perform the configuration step after the restart, not before.)



- If you change the HTTP port on a server where it affects the Maestro User Interface component, then you also have to change the **LUI Access URL for Users** setting in the Administration Hub to reflect this port change.

This may affect the **Default LUI Access URL for Users** under **Global Settings > Maestro User Interface > Default URL Settings** and/or the individual **LUI Access URL for Users** setting on group or user level.

You may also have to change the **LUI Access URL for Admin** setting (on default level) and/or the **Subscriber Access URL** setting (on group or user level), if these settings do not inherit the **LUI Access URL for Users** setting.

- If any LISTSERV Web Interface Links have been defined for the Maestro User Interface (in the Administration Hub), then these must be refreshed so that they are aware of the new HTTP port: In the Administration Hub, select **Global Settings > Maestro User Interface > LISTSERV Web Interface Access > LISTSERV Web Interface Links**. If you have any entries in the table, then you need to refresh each entry as follows (do this after applying all the other steps described here and after you have restarted LISTSERV Maestro):

Click the **Edit** link of the entry. On the next page, take note of the host name that is currently selected in the drop-down list at the top, then copy the access URL from the edit box at the bottom (for later paste). Then, click the **[Delete Link]** button. This deletes the old interface link and you are now back on the previous page. Now, click the **[Create New Link]** button. On the next page, select the same server as was used before from the drop-down list at then top, then paste the access-URL that you copied before into the edit box at the bottom. Click **[OK]** to re-create the link. You have deleted the old interface link and replaced it with an equivalent new one that is now aware of the new HTTP port used by LISTSERV Maestro.



Important: If the LISTSERV Web Interface (WA) is also being hosted by LISTSERV Maestro's Tomcat, then the HTTP port change will also affect the WA, which means that also the access URL which you copied and pasted above needs to be changed (to reflect the new HTTP port). In this case, before you click "OK", first edit the pasted access URL so that it contains the new HTTP port (this is not necessary if the WA is being hosted on a different server).



- If you change the HTTP port on the server that affects the Maestro Tracker component, then you also have to change the **Tracking URL** setting in the Administration Hub to reflect this port change.

This may affect the **Default Tracking URL** setting under **Global Settings > Maestro User Interface > Default URL Settings** and/or the individual **Tracking URL** setting on group or user level.



- If you change the HTTP port on a server where it affects the Administration Hub component, then you also have to change the **HUB Access URL for Users** setting in the Administration Hub to reflect this port change.

This may affect the **Default HUB Access URL for Users** setting under **Global Settings > Maestro User Interface > Default URL Settings** and/or the individual **HUB Access URL for Users** setting on group or user level.

You may also have to change the **HUB Access URL for Admin** setting (on default level), if it does not inherit the **HUB Access URL for Users** setting.

- The Maestro User Interface and the Administration Hub User Interface are both accessed using an HTTP port. This implies that if this port is changed, it will no longer be possible to access these interfaces by entering the default URL into the location field of a browser. Instead, it is necessary to add the port number (with a colon ":") to the URL. For example, if the HTTP port is changed to 8080, then the access URL will need to include the port number as shown below:

```
http://your_host:8080/loi
```

- The shortcuts to access the Maestro User Interface and the Administration Hub User Interface that are installed in the Windows start menu do not include any port information. They expect the user interfaces to be accessible on the standard port 80. If this port is changed, then it is necessary to edit these shortcuts and add “:yourPort” to the URL, as described above.
- Changing the HTTP port also affects the `CompileAll` command (this command is a tool to pre-compile all pages before first use – see the Installation Manual for details). Usually this command is only executed once, right after installation. However, if `CompileAll` needs to be run again (for example after an upgrade installation), and the HTTP access port has been changed for the LISTSERV Maestro installation, the following files must be edited:

```
[maestro_install_folder]/commands/compile/hub.host  
[maestro_install_folder]/commands/compile/loi.host
```

- The “`hub.host`” file is located on the server where the Administration Hub component is installed, while the `loi.host` file is located on the server where the Maestro User Interface component is installed. If both components are on the same server, then the two files will be as well.
- The file can be edited with any text editor. It contains a single line, comprised of the access-URL (including host-name and port) for the Administration Hub and the Maestro User Interface component, respectively. Change it so that it contains the new HTTP port with a colon “:” after the host name (or leave out port and colon if the port is the standard port 80). For example, if the HTTP port was changed to 8888, then the Maestro User Interface entry must look like this:

```
http://yourhost.domain.etc:8888/loi
```

The entry for the Administration Hub will look similar, only with `/hub` at the end.

If the port is changed back to the standard 80, then either include `:80` instead of the `:8888` shown above, or just leave out the port and the colon.

```
http://yourhost.domain.etc/loi
```

- If LISTSERV Maestro is installed behind a firewall (which is advisable) and the Maestro User Interface and/or the Administration Hub User Interface needs to be accessible from a computer outside the firewall, the firewall must be configured to allow access on the configured port instead of the standard HTTP port.
- Similarly, if the Maestro Tracker component is installed behind a firewall, then the firewall must be configured to give all outside users access to the server where Maestro Tracker is installed on the port that is configured for HTTP access. This is normally port 80, but can be a different port if the port was changed as described above.



Important: The whole tracking mechanism of LISTSERV Maestro *will not work* if the Maestro Tracker component is installed behind a firewall in a way such that outside clients do not have access to its configured HTTP port.



Note: If you change any of the access URLs in the HUB as described above (to allow for a non-standard port), and you make a mistake so that now the wrong access URL is submitted, then this may have the effect that users may no longer be able to login to LUI. This may even affect the admin login to LUI and HUB, in which case the admin can not even login normally to correct the mistake. See Section 31 [Emergency Admin HUB Access](#) for how to login in to the HUB anyway in such a case, to be able to correct the mistake.

14.2.2 Configuring the HTTPS Port



To configure the HTTPS port (for SSL), you need to edit the “SecurePort” entry in the file:

```
[maestro_install_folder]/conf/tomcat.ini
```

Example: `SecurePort=5443`

If you leave out the entry from the tomcat.ini file, the default port “443” will be used for HTTPS connections.



Note: If you have several LISTSERV Maestro components installed on the same server, then all of them will be affected by this change. This means that it is **not** possible to use different HTTPS ports for each of the components, if the components are installed on the same server. However, if the components are installed on different servers, then they can use different HTTPS ports.



Important: These changes will only be effective after a restart of the component in question.

When changing the HTTPS port, you need to be aware of the same issues (for example adjusting certain HUB settings), as when changing the HTTP port. Please see Section 14.2.1 [Configuring the HTTP Port](#) about what those issues are (only applied to HTTPS instead of HTTP).

14.2.3 Configuring the Internal Communication Port

This port can be configured independently for each component. However, if the components are installed on the same server, then they must all use the same internal communications port.

- To configure the communication port for the Administration Hub component, edit the following file:

```
[maestro_install_folder]/hub/hub.ini
```

Edit or add the entry `RMIPort`. If the entry is not present or is commented out, the component defaults to port 1099. For example: `RMIPort=5310`

In addition, it is necessary to edit the INI-file of each component that works together with this Administration Hub component. This is usually one Maestro User Interface and one Maestro Tracker component.

- For the Maestro User Interface component, edit the file:

```
[maestro_install_folder]/lui/lui.ini
```

Edit the entry `HubRMIPort` (if this entry is not present or is commented out, it defaults to port 1099).

- For the Maestro Tracker component, edit the file:

```
[maestro_install_folder]/trk/tracker.ini
```

Similarly, in this example, edit the entry `HubRMIPort` in the same way as described above for the Maestro User Interface component.

- To configure the communication port for the Maestro User Interface, edit the following file:

```
[maestro_install_folder]/lui/lui.ini
```

Edit or add the entry `RMIPort`. If the entry is not present or is commented out, the component defaults to port 1099. Example: `RMIPort=5310`

- To configure the communication port for Maestro Tracker, edit the following file:

```
[maestro_install_folder]/trk/tracker.ini
```

Edit or add the entry `RMIPort`. If the entry is not present or is commented out, the component defaults to port 1099. Example: `RMIPort=5310`

- In addition, edit the Maestro User Interface INI file that communicates with the Maestro Tracker component. On the server where the Maestro User Interface is installed, edit the file:

```
[maestro_install_folder]/lui/lui.ini
```

Edit or add the entry `TrackerRMIPort`.

14.2.4 Configuring the Tracker Communications Port

This port is only used by the Maestro User Interface (LUI) to communicate with Maestro Tracker component (TRK). It can easily be configured using the Administration Hub. Simply enter the Administration Hub, click on the **Global Settings** menu, select **Maestro Tracker**, and then select **General Administration**. Edit the port number. Click **[OK]** to save. The change will be effective immediately.

14.2.5 Configuring the Internal Database Connection Port

LISTSERV Maestro comes with an internal database that can be used as the system database. The Internal Database Connection port is only used by the Maestro User Interface component when it is configured to use this internal database. To configure it, edit the following file:

```
[maestro_install_folder]/lui/database/my.ini
```

In this file, find the entry `port` both in the `[client]` and `[mysqld]` sections. Edit the value of *both* of these entries to change the database connection port. In addition, edit the following file:

```
[maestro_install_folder]/lui/lui.ini
```

Edit the “MySQLConnectorJDriverPlugin.databasePort” entry or the “MySQLDriverPlugin.databasePort” entry (whichever is present) to point to the same port number. These changes will only be effective after restarting the Maestro User Interface component.

14.2.6 Configuring the Application Server Shutdown Port

To configure the Server Shutdown port, you need to edit the `ShutdownPort` entry in the Tomcat INI file:

```
[maestro_install_folder]/conf/tomcat.ini
```

Example: `ShutdownPort=90`

If this entry is not configured, then the default port 8007 will be used.

Section 15 Defining the IP Addresses and HTTP Ports



By default, LISTSERV Maestro binds the HTTP-port on all IP-addresses of the server it is running on, and it binds the same HTTP-port on all these addresses. I.e., if the server has several addresses, then a client will be able to access the Maestro User Interface, the Administration Hub and Maestro Tracker (depending on which components are installed) on the same HTTP-port by using any of the server's addresses.

If this default behavior is satisfactory, then no changes to the LISTSERV Maestro configuration need to be made.

However, if for some reasons it is required that LISTSERV Maestro binds only to some of the available IP-addresses on the server, and/or if you want LISTSERV Maestro to bind different ports on different IP-addresses, you need to edit the "BindAddress" entry in the file:

```
[maestro_install_folder]/conf/tomcat.ini
```

Specify one or several comma-separated IP-addresses, optionally with ports. LISTSERV Maestro will then bind only to the specified addresses, using the specified ports. If for an address no port is specified, then the default http port (as configured by the Port setting in the tomcat.ini) will be bound on that address.

Examples:

- Bind only one address, using the default port:

```
BindAddress=192.168.1.1
```

- Bind three different addresses, using the default port on each:

```
BindAddress=192.168.1.15,192.168.1.16,192.168.1.24
```

- Bind three different addresses, using the default port on the first, but different ports on the other two:

```
BindAddress=192.168.1.15,192.168.1.16:8081,192.168.1.24:8082
```



Note: If you have several LISTSERV Maestro components installed on the same server, then all of them will be affected by this change. E.g., it is not possible to use different bindings for each of the components, if the components are installed on the same server. However, if the components are installed on different servers, they can use different bindings.



Important: This change will only be effective after a restart of the component in question.

Section 16 Installing Behind a Firewall

Any network that is connected to the Internet is usually protected by some form of firewall, often in conjunction with different kinds of “demilitarized zones” and other security measures. If there is a desire to install the components of LISTSERV Maestro behind a firewall, or in different protection zones so that some are behind and others are in front of the firewall, it is necessary to take into account the communication channels between the separate components.

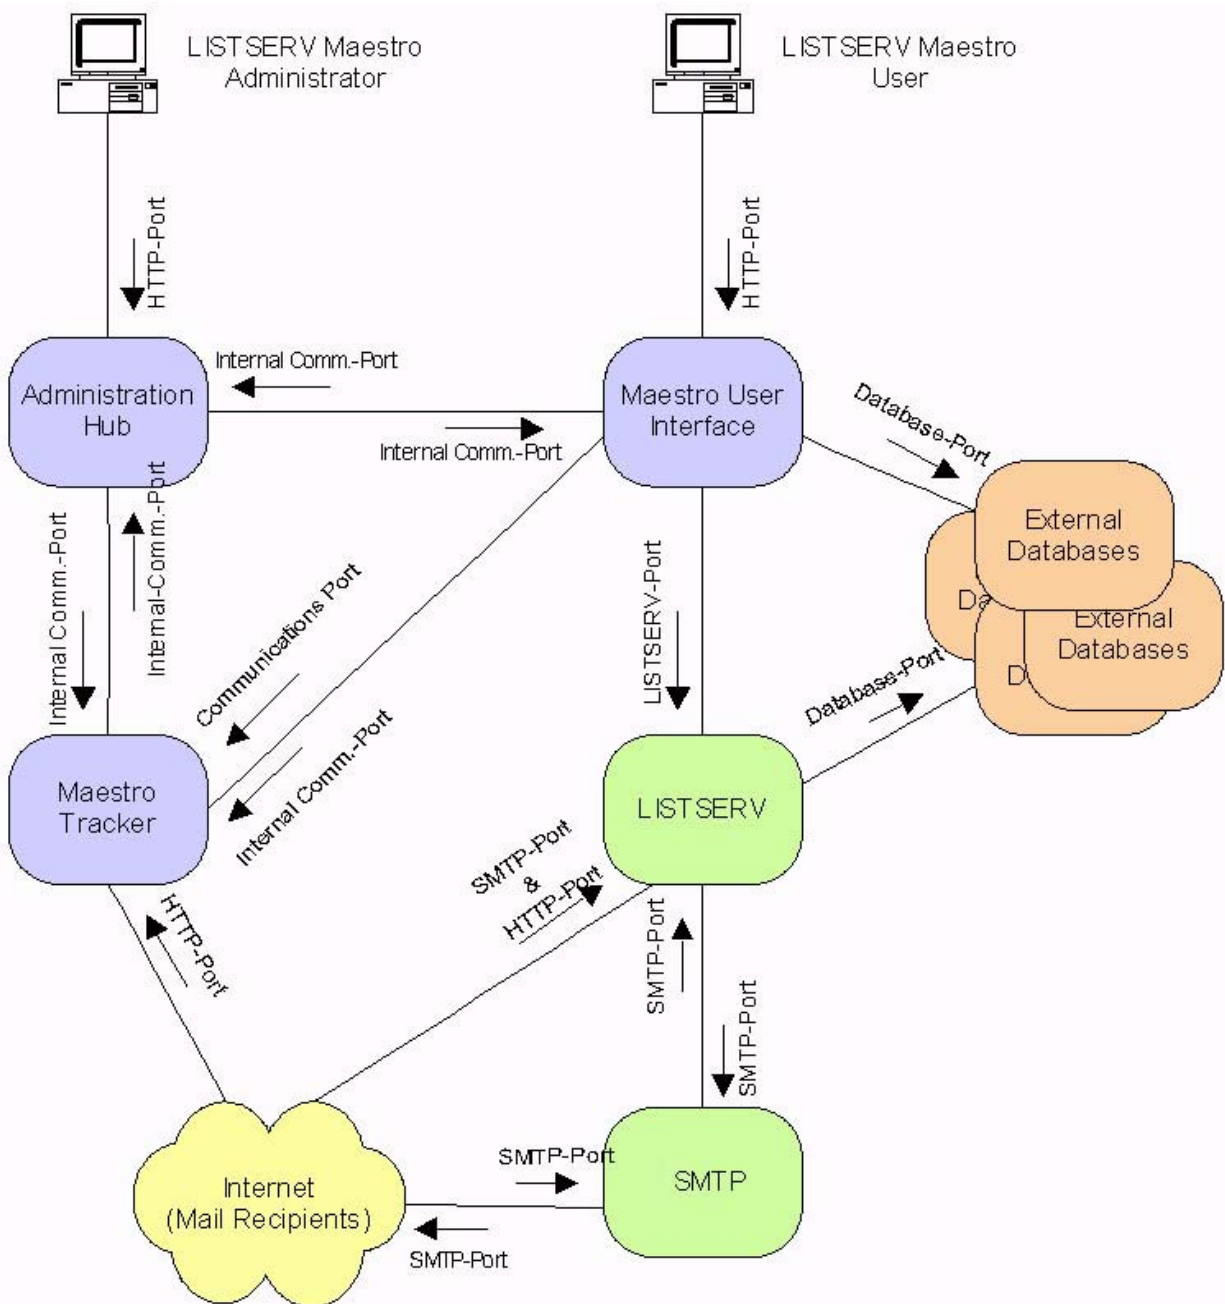
Communication occurs exclusively using TCP ports (see the Section 14 [Using Non-Standard Ports](#) for more information). If the components are installed behind, in front of, or on both sides of a firewall, then the firewall needs to be configured to let communication through on certain ports between certain servers.

At each communication line, a labeled arrow illustrates the direction of the communication between the two components, and the port used for this communication. The communication can go in one direction or both directions. However, if the communication goes in both directions, then an open port is required on both sides.

The port label definitions are:

- **HTTP Port** – Used for standard HTTP access, using a web browser. This is also used to transfer the tracking events from the Internet (from the email messages that were sent) to the Maestro Tracker component. The standard HTTP Port is **80**.
- If HTTPS access to the Administration Hub and/or the Maestro User Interface component is being used, then the HTTP Port from the Maestro Administrator to the Administration Hub and/or the HTTP Port from the Maestro User to the Maestro User Interface should be substituted with the **HTTPS Port**, for which the standard is **443**.
- **SMTP Port** – Used for standard SMTP communication, during the sending and receiving of email. The standard SMTP Port is **25**.
- **Internal Communication Port** – Used for communication between the separate LISTSERV Maestro components and the Administration Hub. The standard Internal Communication Port is **1099**.
- **Communications Port** – Used for special communication between the Maestro User Interface and the Maestro Tracker component to transfer tracking events to the Maestro User Interface component (for reports). The standard Communication Port is **7000**.
- **LISTSERV Port** – Used by the Maestro User Interface component to access the external LISTSERV component. The standard LISTSERV Port is **2306**.
- **Database Port** – Used by the Maestro User Interface component to access the external database component. The standard Database Port depends on the database used.

Figure 16-1 Component Communication Pathways



All the components shown in the figure (except for the Internet, LISTSERV Maestro Administrator, and LISTSERV Maestro User) may reside on a single server or may be distributed over different servers, up to the maximum distribution of a dedicated server for each of the components shown (or multiple servers in the cases of LISTSERV and SMTP).

When two components are installed on the same server, a firewall will not stop the communication between the two (except if the firewall is installed on the same server, where the firewall may close the ports the components use to communicate). However, if some components are installed on separate servers, a firewall may sit between the two.

Most commonly, a firewall will separate the Internet from the other components. The other components may also be installed in a way that has a firewall between them.

Imagine the firewall as sitting “on top” of the connection between two components.

If that is the case, then the firewall must be configured so that it allows communication between the two components, as specified by the arrow(s) associated with the connection the firewall guards. The direction of the arrow shows the direction the port should be opened, and the label of the arrow defines which port needs to be open.

For most components, the safest method will be to open the firewall for only the required port(s) in the required direction(s), and between the IP addresses of the servers where the components reside.

For example, if there is a firewall between the Maestro Tracker and the Maestro User Interface component, open the Communications Port and the Internal Communications Port *only* in the direction from the Maestro User Interface host to the Maestro Tracker host. Open both ports *only* for the IP address involved. This limits the possible security breaches in the case of an unauthorized person gaining access to one of the component servers.

There are some exceptions:

- If there is a firewall that separates the Internet from the other components (as is advisable), open the HTTP and SMTP ports from the Internet to the respective components as shown in the diagram, and open them for all incoming IP addresses, not just for a specific one. In addition, it is necessary to open the SMTP port for outgoing communication originating from the LISTSERV and SMTP servers.
- Similarly, if there is a firewall separating the Internet from the other components as described above, and both the Maestro Administrator and the Maestro User need to be able to connect to LISTSERV Maestro from the Internet as well as the local intranet behind the firewall, then the HTTP port to the Administration Hub and Maestro User Interface components for all incoming IP-addresses must also open. In this case, LISTSERV Maestro’s login security will be relied upon to disallow unauthorized access to these components.

Allowing the Application Server Shutdown Port, (default 8007) access through the firewall is not a concern, as this port is only ever used locally for communication between two processes on the same server. If there is a firewall on the server itself, this port might also have to be opened. Simply check if the L-Soft Tomcat server still reacts to the `stop` command. If not, then the port needs to be opened.

Section 17 Restricting Access to Components

The administrator can restrict access to LISTSERV Maestro in two different ways. The first way of restricting access is based on the IP address of the computer (where the browser is running) that is used to access the component. The second way of restricting access is to disallow concurrent access with the same user account. This will limit users from logging in twice with the same user account at the same time.

17.1 IP Address Restrictions

Each of the LISTSERV Maestro components (Administration Hub, Maestro User Interface, Maestro Tracker, and the subscriber access pages for hosted datasets) can be configured to restrict access based upon the IP address of the computer that is used to access the component where the browser/email-client is running. This means that it is possible, for example, to define that everyone (all IP addresses) is allowed to access the Maestro Tracker component, but only certain IP addresses (a local subnet, perhaps) are allowed to access the Maestro User Interface and Administration Hub components. If access is not allowed for a certain address, then a client from that address will receive a 403: `Forbidden` error when attempting to access the restricted component.

By default, no component access restrictions are in effect. To add access restrictions, it is necessary to add a new `Restrict.CONTEXT.ID` entry into the Tomcat INI file:

```
[maestro_install_folder]/conf/tomcat.ini
```

Each such entry must look something like this:

```
Restrict.CONTEXT.ID=NETWORK/MASK
```

with the following replacements:

- **CONTEXT:** Replace with the context name for which you want to introduce a restriction. Usually you will probably want to restrict access to the Maestro User Interface and/or the Administration Hub, for which the matching context names are “lui” or “hub” respectively.

Other possible context names are “trk” and “list” (for Maestro Tracker and for the subscriber pages of hosted lists, although it usually does not make sense to restrict access for these contexts), and “archive” and “scripts” (these two being contexts used by the LISTSERV user interface “WA”).

- **ID:** Replace with any ID-string that uniquely identifies the “Restrict” entry from all other “Restrict” entries in the same context. Which kind of ID-string you use is up to you, but you should limit yourself to alpha-numeric characters and make sure that you do not use the same ID-string for two “Restrict” entries with the same context name (i.e. two “Restrict” entries must at least differ in their “CONTEXT” or in their “ID” value, there must never be two entries where both “CONTEXT” and “ID” are the same).
- **NETWORK:** Replace with the dot-separated IP-address of the subnet to which you want to grant access to the given context (like “192.168.1.0”).

- **MASK:** Replace with the dot-separated subnet-mask for the subnet specified above (like “255.255.255.0”).

It is important to understand that the listed IP-address ranges or addresses are the addresses which are **granted** access. All unlisted addresses are thus implicitly **denied** access to this context.

If no such restriction entry is present for a certain context at all, then access to this context is **unrestricted** (this is the default for all contexts after installation).

In other words: If for a context there is no entry at all, then access to that context is unrestricted. If there is at least one entry, then access to that context is restricted and access is allowed only for the addresses listed in the entry (or entries) of that context.



Important: Because of the way the Maestro Tracker functions (by accepting tracking events from mails sent all over the internet), the Maestro Tracker component must be accessible to everyone, i.e. you should not specify any restriction entry for the “`trk`” context. For similar reasons, you should also not specify any restriction entry for the “`list`” context, so that everyone has access to the subscriber pages of the hosted datasets (unless you have a policy to restrict access to these pages, for example if you are using them only for internal purposes).

After you have saved the modified `tomcat.ini`, you need to stop and restart LISTSERV Maestro to make it aware of the changes.

If you have distributed the components of LISTSERV Maestro to several servers, then you might need to edit the `tomcat.ini` file of several of these servers, depending on which components you want to restrict.

For example, if all three, the Administration Hub, the Maestro User Interface and Maestro Tracker are installed on separate servers, then you would typically not add a restriction entry on the Maestro Tracker server (since Maestro Tracker needs to be accessible to all), but you might want to add restriction entries both to the `tomcat.ini` of the Administration Hub (using the “`hub`” context) and of the Maestro User Interface server (using the “`lui`” context).

Examples:

```
Restrict.lui.0=192.168.1.0/255.255.255.0
```

This would restrict access to the Maestro User Interface (“`lui`”) and only allow access for computers in the subnet range 192.168.1.0 through 192.168.1.255. Computers with any other IP address would not be allowed to access the Maestro User Interface.

Access to all other components (for example the Administration Hub, Maestro Tracker or the subscriber pages of hosted lists) would remain unrestricted.

```
Restrict.lui.0=192.168.1.0/255.255.255.0
```

```
Restrict.lui.1=192.168.6.21/255.255.255.255
```

```
Restrict.hub.0=192.168.6.21/255.255.255.255
```

This would restrict access to the Maestro User Interface (“`lui`”) and only allow access for computers in the same subnet range as above and additionally also for the single computer with the address 192.168.6.21.

Also, access to the Administration Hub (“`hub`”) is restricted and access is allowed only for this one same computer with address 192.168.6.21.

Access to the Maestro Tracker, the subscriber pages of hosted lists, and the LISTSERV Web interface (if being served by LISTSERV Maestro) remains unrestricted.

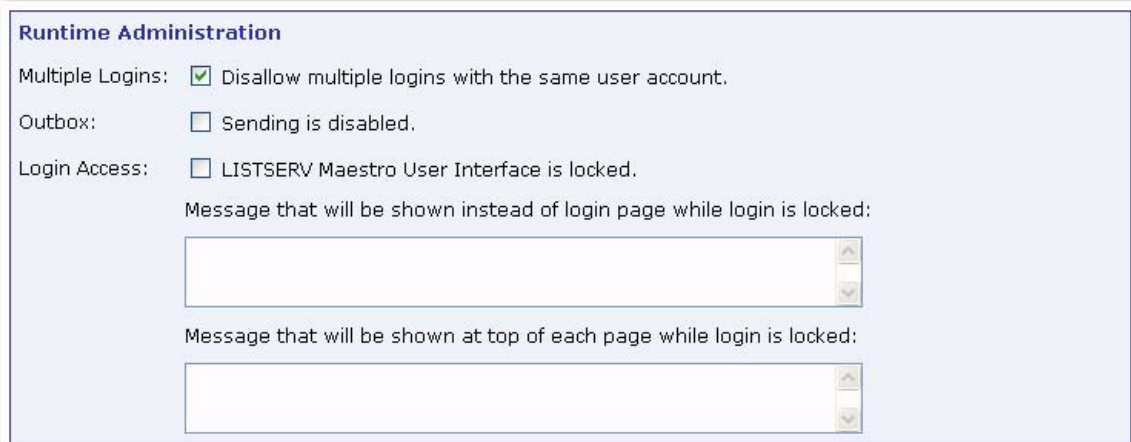
17.2 Disallowing Concurrent Access with the Same User Account

If there is an organizational reason or policy that dictates this restriction, the administrator has the option of allowing or disallowing users to log in twice with the same user account at the same time. The default setting does allow concurrent access. This restriction should only be used in special cases and with an understanding of the problems associated with using it.

There is usually no reason to disallow concurrent access to LISTSERV Maestro. If two users are logged in with the same account from different workstations, Maestro handles each login session separately; therefore, the two sessions will not interfere with one another. If multiple users need to access and manipulate the same data, then it is generally a better idea to assign separate accounts in the same group to each user rather than allowing them to share a single account. Doing so not only eliminates the need to disallow concurrent access, but it also allows for more detailed logging (i.e. log files that show which user performed particular actions).

To change the default to disallow concurrent access, click on the **Global Settings** menu, select **Maestro User Interface**, and then select **General Administration**. The General Administration of Maestro User Interface screen opens. In the **Runtime Administration** section, check **Disallow multiple logins with the same user account**.

Figure 17-1 Multiple Logins



Runtime Administration

Multiple Logins: Disallow multiple logins with the same user account.

Outbox: Sending is disabled.

Login Access: LISTSERV Maestro User Interface is locked.

Message that will be shown instead of login page while login is locked:

Message that will be shown at top of each page while login is locked:



Important: Disallowing concurrent access with the same account is not recommended. If it is necessary for some reason, please pay attention to the warnings issued below about potential problems associated with the use of this feature.

Disallowing concurrent access will affect the behavior of the Maestro User Interface. If a user logs in with a certain account, and another user is already logged in with the same account, the system will not accept the second login right away, but will instead do the following:

- If the second login attempt comes from a different workstation, the user attempting the second login is given the message “Logon failed: Someone is already logged in

with the given account from a different workstation. Please use a different account for login.” The user is not logged in. However, the user may still use a different account that is not currently in use to log in.

- If the second login attempt comes from the same workstation, the user is informed that a previous session is already active from the same workstation. The user is then asked whether to cancel the second login, or proceed with the second login and log out of the previous session. If the user cancels the second login, the previous session will be unaffected, but the second login attempt will fail. If the user does not cancel the second login, the previous session will be logged out and the second session will log in.

A second login attempt from the same workstation may happen in situations similar to these:

- A user has one browser window open, in which the first login session is active. The user opens a second window and tries to log in again with the same account. In this case, the user will be notified that there still is a session open from the workstation and that proceeding with the second login will log out that first session. Most users will probably cancel the second login instead and continue using the first session.
- A user has been using a first login session in a browser and has closed the browser without logging out properly. Since the system has no way of knowing that the user has closed the browser, it will still keep the user’s login session active. Since the browser is closed already, the user has no way of “going back” to that session to log out properly.
- This is usually not a problem, since the system will log out the session automatically after a certain timeout period has passed (usually 90 minutes). However, if in the meantime the user opens a new browser window and tries to log in again with the same account, the user will be notified that there is already a session logged in from the workstation, and that proceeding with the second login will automatically log out that first session. Since the first session is the one that the user no longer has access to, the user will proceed with the second login.

LISTSERV Maestro makes the determination of whether a second login attempt comes from the same or from a different workstation by looking at the IP address of the workstation used to make that attempt.

This approach has some caveats of which to be aware (illustrated in the scenarios below).

Problem Scenario #1: NAT Access

If a group of users is accessing the Maestro User Interface using a local subnet with local addresses, and a router with NAT (Network Address Translation) or some other method of address mapping is used to connect to the Internet, and the Maestro User Interface is on the “other” side of that router, then to the Maestro User Interface, all users will appear to be using the same workstation, since they will all have the same IP address, namely that of the router.

In this case, the Maestro User Interface will handle all login attempts as if they were originating from the same workstation, which may result in the following confusing or

even harmful situation. One user is logged in with an account from workstation A. Now another user tries to log in with the same account, only from workstation B. Both workstations will appear to the Maestro User Interface as one and the same, since both will be using the same IP address externally. The result is that the second user will be notified that there is another session already active from the workstation with the same account. The user will have the option of proceeding with the login and canceling the previous login. This other session would in fact be the session of the first user and by logging in, the second user would log out the first user, disrupting the workflow.

To work around this situation, make sure that all users are using different accounts, and that the passwords are kept secret, so that no other user can use a colleague's account to log in from a different computer and thus log out that colleague.

Problem Scenario #2: Dial-Up Access

If a user is connected to the Internet with a dial-up modem connection as provided by most ISPs, the workstation's IP address is usually assigned dynamically each time the user connects, meaning a different IP address will be assigned each time a connection is made. This may cause the following situation to happen:

The user opens a browser and logs into the Maestro User Interface with a certain account. The user then closes the browser without logging out properly, so that the session will continue to be active until the timeout has expired. The user then disconnects the Internet connection. Shortly thereafter, the user reconnects to the Internet, opens another browser, and tries to log in with the same account. This time, the user is very likely to be assigned a different IP address from the previous connection. The Maestro User Interface will interpret this as a different workstation logging in to the same account. As a result, the Maestro User Interface will report that the account is currently in use from a different workstation and will not accept a login with that account.

The user now has no choice but to wait for the 90 minutes timeout to expire, before logging in again with the same account. To cancel the previous login, the user would have to access the Maestro User Interface using the same IP address as before, which is extremely unlikely with this kind of dynamic address assignment. To avoid this problem, the user should always remember to log out properly. If the browser is closed accidentally without logging out, but before the modem is disconnected, a new browser session should be opened so that the user can log in again, canceling the previous session, and then log out properly.

To moderate this problem, the administrator may configure the session timeout of the Maestro User Interface to be shorter than the default of 90 minutes, so that in the worst case, the user does not have to wait as long to log back in.

The timeout for the Maestro User Interface is configured in the following file:

```
\Program Files\L-Soft\webapps\lui\WEB-INF\web.xml
```

Example

```
<!-- 1.5 hrs session timeout -->
<session-config>
    <session-timeout>90</session-timeout>
</session-config>
```

The value of 90 determines the session timeout in minutes. Set it to a suitable value, save the file and restart the Maestro User Interface.

The same setting can be changed for the Administration Hub by editing the file

```
\Program Files\L-Soft\webapps\hub\WEB-INF\web.xml
```

17.3 Securing Access Against Dictionary Attacks

A dictionary attack is a technique to gain illegal access to a system by employing a list of words in a dictionary automatically to determine the login password for a given user account. The effectiveness of such an attack can be reduced by only allowing a limited number of invalid login attempts and by locking access to the account for a certain time. (Locking means that the login is denied even if the correct password is supplied.) LISTSERV Maestro supports this form of login locking in the Administration Hub and in the LISTSERV Maestro User Interface component.

17.3.1 Securing the Administration Hub

To secure the administrator's account of the Administration Hub, click on the **Global Settings** menu, select **Administration Hub**, and then finally select **General Administration**. In the **Advanced Security Options** section, enter the maximum number of unsuccessful login attempts and the login locking time in minutes.

If the administrator's account is already locked due to too many login attempts and the configured login locking time is very long, supply the following value in the hub.ini file:

```
UnlockLockedAccess=true
```

Then, retry to login as administrator's with the correct password. Login access is enabled again and the entry from the hub.ini file has been removed. If a system restart is an option (e.g. because currently no important mail job deliveries are being processed), then restart the system to unlock access again.

17.3.2 Securing the LISTSERV Maestro User Interface

To secure all accounts of the LISTSERV Maestro User Interface, click on the **Global Settings** menu, select **Maestro User Interface**, and then finally select **General Administration**. In the **Advanced Security Options** section, enter the maximum number of unsuccessful login attempts and the login locking time in minutes.

If any LISTSERV Maestro User Interface account is already locked due to too many login attempts, click the **[Unlock all currently locked accounts]** button. If a system restart is an option (e.g. because currently no important mail job deliveries are being processed), then restart the system to unlock all locked accounts again.

Section 18 Securing Access with SSL

As described in the section regarding the use of non-standard ports and installing behind a firewall, the LISTSERV Maestro components can be configured in a way that users and/or administrators can access the Maestro User Interface and/or the Administration Hub with a web browser from anywhere on the Internet. This feature allows, for example, LISTSERV Maestro to be set up in an ASP-environment, where the customers access the application over the Internet.

Providing access to users from the Internet exposes the risk of unauthorized access. LISTSERV Maestro uses password authentication as a first barrier against intruders. However, network traffic is a public affair. Anyone with the right knowledge and access to certain nodes in the network may eavesdrop on the communication between the user's browser and the LISTSERV Maestro server. Intruders may gain knowledge about the data that is sent to the user's browser (for display) and sent back to the server (to trigger a certain action or to submit settings the user made). Even more dangerous, the intruder could find out the user name and password that the user or administrator employs for login, and could then log him/herself in with the same account.

If security is a concern, consider securing access to the LISTSERV Maestro servers with encrypted communication, so that intruders cannot listen in on the communication between browser and server, and cannot gain knowledge about the data exchanged or spy out passwords. All standard browsers support encrypted communication using Secure Sockets Layer (SSL), and the HTTPS protocol to access web pages, instead of the normal HTTP protocol.

The following subchapters describe how you can configure your LISTSERV Maestro server(s) to use SSL. Since topics such as encryption, server certificates, and trusted authorities are so complex, an introduction is presented in Section 18.1 [Introduction to Secure Communication](#) to assist understanding the concepts involved, making the execution of the required steps easier. Implementation instructions start in Section 18.2 [Which Components Should Be Secured?](#).

Securing access with SSL (HTTPS) as described in this section is a separate issue from authenticating and encrypting communication between the components of LISTSERV Maestro, even though the two have many similarities and can even be combined.



Note: To authenticate and encrypt the communication between the separate components of LISTSERV Maestro please refer to the L-Soft White Paper entitled [Authenticating and Encrypting Communication between LISTSERV Maestro Components](#).

18.1 Introduction to Secure Communication

This section is intended to provide a short introduction about the basics of secure communication. Please see the many publications about this topic for more details.

For successful encrypted communication to take place, one partner holding an "encryption key" encrypts the data. The data is transferred to the second partner and decrypted using the same "key". One requirement is that both communication partners

know the encryption key so that the receiving partner may decrypt the data that was encrypted by the sending partner.

With online communication, however, this is more complex. Both partners (the browser and the server) are most likely communicating with each other for the first time, and do not have a common encryption key that is known only to them. So, when the connection is first established, the two partners must secretly decide at the spur of the moment on an encryption key that will be used for the rest of the communication (this is a simplified view of the matter, but it explains the basics).

Assuming that both partners have decided what key to use, they can now communicate in an encrypted manner. There is still the problem of being sure that each partner is actually communicating with the partner they think they are communicating with. An analogy to this problem can be found in real life. Suppose that two employees of two partner companies meet in a hotel to exchange confidential information. The two have never met each other, but they know each other's names and home addresses. How can each of them be sure that the other person they meet in the lobby of the hotel is actually the person they are supposed to meet and not an impostor?

An impostor could act as a "man-in-the-middle". He meets with employee A of corporation A-Corp in the lobby and poses as employee B of corporation B-Corp. Thus, he gains confidential information from A and goes into the bar where he meets the real employee B. Here he poses as employee A from A-Corp, gives the confidential information from the "real" A to B and receives similar information back from B. Finally, he goes back into the lobby relays the information he received from B to the "real" employee A. On his way from the lobby to the bar and back, he made copies of the information he was carrying. In the end, both employee A and B are unaware that they did not talk to their "real" counterparts, but to an impostor that acted as a "man-in-the-middle", and the impostor goes back to his employer C-Corp with the confidential information he gained from their competitors.

On a network, this "man-in-the-middle" attack is even easier to mount. The only thing that a server and a client know of each other is their network addresses, which can easily be forged. In real life, the two employees of A-Corp and B-Corp would probably request to see some picture ID with name and home address of their communication partner. They would then compare the picture on the ID with the person they are talking to and verify that the name and address on the ID matches the ones they have previously been told. If the ID matches the person, they would be confident that they are talking to the right person.

However, in doing so, they actually implicitly trust a third party that has not yet been involved. This would be the agency that issued the picture ID. By accepting the ID, they trust that the agency has created an ID that is hard to forge. They also trust that this agency has, in turn, verified that the person they issued the ID to really is the person he claims to be. If employee A had tried to use his library card for identification, then employee B would probably have rejected it as improper identification, because she would not trust either that the clerk in the library responsible for issuing the ID really did a thorough check of A's identity, or she did not trust the security features of the ID, (these days anyone can create an authentic looking ID with the help of a color laser printer). Instead, she would probably request a "proper" ID like a passport or driver's license.

With online communication, the problem of identifying the communication partner is solved very similarly. In the online world, "certificates" fulfill the role of a "picture ID" in

real life. A certificate asserts that the owner of the certificate is, in fact, the entity it claims to be. For example, a certificate could assert that the server with the host name "host.somecorp.com" actually is a server that belongs to SomeCorp, and that it is not an impostor's server.

How can such credibility problems be solved? Simply falsifying a file that states, "Yes, the server 'host.somecorp.com' is indeed a server belonging to the SomeCorp Corporation," would not be cumbersome for an imposter. To guard against this, the certificate is digitally signed by a trustworthy authority, so that it now reads, "Yes, the server 'host.somecorp.com' is a server belonging to the SomeCorp corporation, and we, the people from TrustCorp have verified that this is indeed so."

The digital signature is very useful because it prevents anyone from tampering with the certificate. If even a single letter (or byte) in the text of the certificate is changed, the signature will no longer match and the certificate will be invalid. However, a last problem remains: how to test the validity of the signature? The digital signature of the certificate was created using a signature key. The signature key consists of two parts, a private key and a public key. The signer (the signing authority) uses both parts to create the signature. The private key is held secret by the signer so that only the signer is able to use it to create a signature and a signature cannot be created with the public key alone.

The public key on the other hand is made public. It can be used by anyone who desires to test the validity of a digital signature that is supposed to come from the owner of the public key. With a certain algorithm, the signature is tested against the public key, producing a result that states (if the signature was valid): "Yes, the data signed by this signature has not been tampered with and it was signed with a private/public key pair, where the public key matches the public key that was used to test the signature."

If the signature was not valid, the result could be, "The data was tampered with since it was signed," in which case the data seen by the recipient is not the original data that the signer saw. Or, the result could be, "The data was signed with a private/public key pair, where the public key does not match the public key used to test the signature." In this case the data was signed by someone else than the owner of the public key the recipient has. In both cases, the signature is invalid.

It is not possible to use any public key that is found anywhere (or given by any one), because who would then guarantee that the public key received really is a key from the entity it supposedly belongs? The origin of the public key that is used to verify the digital signature of a certificate has to be very reliable; otherwise, the "man-in-the-middle" would still have a chance to spy. The intruder would create his own public/private keys with a forged name of "TrustCorp" and his own certificate with a forged host name of SomeCorp. Then he would use his own private key to sign the certificate and would give others the public key claiming, "This is the public key of TrustCorp". If this public key were used to check the validity of the forged certificate, a match would be made leading to the belief that the forged certificate is legitimate. As a result, the attacker would receive the communication and not the server of TrustCorp.

To verify a public key of the signing authority, most web browsers, like Microsoft® Internet Explorer, are already equipped with a list of trusted so called "root certificates." It is not necessary to verify that these certificates indeed come from the entity they claim to, because the browser vendor has already verified this.

The full trust-chain when a browser is used to access a secured is described as follows:

- The browser vendor receives root certificates from the signing authorities, verifying their validity.
- The browser vendor trusts that the root certificates are genuine.
- The browser trusts any signed certificate with a genuine certificate traceable to one of the trusted root certificates. This can be a very short chain such as, for example, “a certificate signed with a root certificate” or a long chain such as for example, “a certificate signed with a certificate that was signed with a certificate ... etc. ... that was signed with a root certificate.”
- The browser trusts any server that has a browser-trusted certificate.

In the real life example, employees A and B both needed picture IDs to verify each other. With online communication, this verification is often only one-sided. For most purposes it is enough that the client is certain about the server it communicates with. It is usually not required that the server is also certain about the client. Therefore, usually only the server has a certificate (which is, down the trust chain, signed by a trusted root certificate) and the client does not.

There are also real-world examples of this. If a car were purchased privately from its former owner, the buyer would most likely request to see a picture ID of the owner during the transaction. Otherwise she would risk unknowingly buying a stolen car. On the other hand, it is not a requirement for the former owner to see the buyer’s ID to check the car’s legitimacy.

To summarize the concepts introduced:

- **Server Certificate** – This certificate asserts that a certain server (with the given host name) actually belongs to a certain organization, so that the server can be trusted and confidential data can be safely communicated. This certificate is digitally signed to prevent tampering and falsification.
- **Trusted Root Certificate** – The trusted root certificate is used to sign the actual server certificate (or another certificate down the trust-chain is used to sign the actual server certificate). Usually the fact that a root certificate is installed together with trusted software (like the browser) already makes it a trusted root certificate. A root certificate can be received by other means (by email, for example). In that case, first verify the certificate before it is rated as “trusted.” To do this, compare the fingerprints on both the sent and received certificates.
- **Encrypted Communication** – This is made possible with the help of an encryption key, which is secretly generated when the communication first begins. Verifying that there is no “man-in-the-middle” while negotiating the encryption key is achieved by verifying the communication partner’s certificate and matching its digital signature, to one of the trusted root certificates, further down the trust-chain.

18.2 Which Components Should Be Secured?

Normally, you should only secure the Maestro User Interface and/or the Administration Hub component with SSL, not the Maestro Tracker component. Because securing the

Maestro Tracker component has several drawbacks, HTTPS is slower and takes more system resources than HTTP, so if you have a high tracking event volume, then using HTTPS puts a higher workload on the Tracker server. Also not all clients may be able to access HTTPS-URLs correctly, which could result in tracking events not being counted and recipients not being redirected to the intended target pages (in case of click tracking). Before you decide to also secure Maestro Tracker, you should take these possible drawbacks into account and evaluate, if they are worth the additional security gain in your situation.



Additional preparations need to be made if you have several components installed on the same server, but only want to secure some of them. For example, if you have all three components on the same server, but only want to secure LUI and HUB, but not TRK (which is recommended, as explained above). In such a case, you need to have at least two separate IP-addresses for the server, each with a separate host name. You would then use one host name and IP-address pair for the components that are to be secured (LUI/HUB) and the other pair for the components that are not to be secured (TRK).

18.3 Obtaining and Installing a Server Certificate

To enable LISTSERV Maestro to use HTTPS by way of SSL, obtain a signed server certificate for the server to be secured. It is not possible to simply obtain any server certificate and use it on any server. The certificate is always bound to the explicit server name that was chosen when the certificate was created. If the LISTSERV Maestro component is moved to a different server (with a different name), or the server is renamed, then a new certificate for the new name would have to be obtained.

Obtaining a server certificate involves three basic steps:

- Create an unsigned certificate with the name of the server being secured.
- Create a certificate-signing request (CSR) from that certificate and send it to a certification authority (CA). The CA first verifies that the requester is genuine, and then returns a signed version of the certificate to him/her.
- Replace the unsigned certificate with the signed certificate returned by the CA.

Certificate administration happens with a command line tool called “`keytool`”, that is installed together with Java. For more information about this tool, and further discussion about certificates and secure communication, see the relevant documentation at Sun’s website:

<http://java.sun.com/j2se/1.5.0/docs/tooldocs/windows/keytool.html>

and

<http://java.sun.com/j2se/1.5.0/docs/tooldocs/solaris/keytool.html>

18.3.1 Securing the Trusted Root Certificate Keystore

As a first step when starting to use certificates, be sure to secure the default keystore for trusted root certificates that is shipped with Java. The Java version that is installed together with LISTSERV Maestro includes a keystore that already contains trusted root

certificates from some CAs (for example VeriSign and Thawte) This keystore is initially protected with the default password “changeit”, which should be changed as soon as possible after the installation of LISTSERV Maestro.

To change the password of the default keystore, execute the following command:

```
[maestro_install_folder]/java/bin/keytool -storepasswd  
-keystore [maestro_install_folder]/java/lib/security/cacerts
```

You will be queried first for the old password (which is `changeit` if it has not been changed since installation of LISTSERV Maestro), and then twice for the new password. Enter a new password with at least six characters, but longer and complex passwords are safer.

18.3.2 Creating an Unsigned Server Certificate

In Java, all certificates are stored in a “keystore,” which is usually a special file protected with a password.

To add a certificate to a keystore, execute the following command:

```
[maestro_install_folder]/java/bin/keytool -genkey  
-alias NAME -validity DAYS -keystore KEYFILE -keyalg RSA
```

with the following replacements:

NAME: The name of the certificate. Can be any name that is not already in use in the keystore file specified (see below), but choose an informative name that helps in recognition of the certificate later.

DAYS: Limits the validity of the certificate. The certificate will expire so many days after the day it was created. This can be any number of days. Usually, when the signing service from the CA is purchased, only a limited period during which the certificate shall be valid is paid for. Choose a number of days for this parameter, which is no shorter than the period purchased from the CA (a little padding here is probably a good idea, to be on the safe side). It is also possible to create a certificate that has a very long validity period (several years), if desired.

KEYFILE: The keystore file to which the certificate will be added. This can be a relative or a full path name. If the file does not exist, it is created. If it already exists, a certificate with the given “NAME” is added to it.

Choose a suitable location and file name for the keystore file that takes into account the special security considerations for this file as outlined below.

Be very careful with the keystore file into which the certificate has been created. Protect this file in two respects:

- Do not lose or accidentally delete this file, as it contains the certificates. New certificates would have to be purchased in this event. Keep a backup at a safe location.
- Protect the file against unauthorized access. Even though the file is password protected, passwords can always be cracked, and an attacker could thus gain access to the certificates.

The tool will first prompt for the entry of the password with which the keystore is protected. If an existing keystore is being used, enter its password. If a filename of a keystore that does not yet exist is given, then a new keystore will be created and it will be protected with the password that was entered at the first prompt (choose a password with at least six characters, remembering that longer and more complex passwords are safer).

Next, the tool will prompt for the following information values. Press RETURN each time to simply accept the default value "Unknown". However, some values must be entered for the certificate to work and some CAs require other values to be filled out. So it is generally a good idea to fill out all values with whatever fits best in each case (see below):

"What is your first and last name?"

Here, the *host name* of the server to be secured with the certificate being created must be entered. Yes, even though the question reads "your first and last name," it is necessary to enter the host name of the computer instead! This should be the same host name that will be used in the URLs to access the server. For example, if the URL is "http://maestro.mycorp.com/lu1", then enter the host name "maestro.mycorp.com" (without the quotes).

"What is the name of your organizational unit?"

"What is the name of your organization?"

"What is the name of your City or Locality?"

"What is the name of your State or Province?"

"What is the two-letter country code for this unit?"

Use the two-letter code that fits the country where the server is deployed, such as US, DE, SE, CH, and so on.

After the last question is answered, a summary of the input and a request for confirmation will appear. Type "yes" and RETURN to accept the input, or "no" and RETURN (or simply RETURN) to reject it (in this case enter the values again until they are satisfactory). After the input is confirmed, the tool takes a few seconds to generate the certificate. When it is done, enter a password at the prompt to protect the certificate. Although generally any password is usable, for the certificate to be usable with LISTSERV Maestro, the same password chosen for the keystore itself must be used. To do so, simply press RETURN without entering anything, so that the default is accepted.

At this point, the certificate has been created, but it is yet unsigned.

18.3.3 Performing a Certificate Signing Request (CSR)

Once an unsigned certificate has been created, generate a certificate-signing request (CSR) from it, which can then be submitted to a certification authority (CA), for example VeriSign.

To generate a CSR for a certificate in the keystore, execute the following command:

```
[maestro_install_folder]/java/bin/keytool -certreq  
-alias NAME -file OUTFILE -keystore KEYFILE
```

with the following replacements:

NAME: The name of the certificate. This must be the name of the certificate that the CSR is created for (the same name that was specified when the certificate was created).

OUTFILE: The file into which the CSR will be written. If the file already exists, it will be replaced with the new file. This can be a relative or a full path name.

KEYFILE: The keystore file in which the certificate is stored. This can be a relative or a full path name.

The command will request the password of the keystore. After it is entered, the file specified as “*OUTFILE*” will be written. This file is a text file that contains the CSR in Base64-encoded form.

Figure 18-1 Example of Base64 Encoded Outfile

```
-----BEGIN NEW CERTIFICATE REQUEST-----
tPnJhsLOuocsbYAmyM1lqiZ5BEVWAnJfZ6kyN/Xft5NFxGIy9Uynz5kODfBwFUgiu98iQKWyMKC/
bGFuZ2VuMQ8wDQYDVQKKEwZMLVNVZnQxEDA0BgNVBAsTB1Vua25vd24xDzANBgNVBAMTBnRlcHBp
6E7Zyl9wkPyVpnlqbnbtXQGAablJInE9/LruaJ1NX1f/NVJgL4vPiDKsU4laGvJHBNhdj+F0uVb
3SIb3DQEBBAUAA4GBAB6XqdfJvhy7dThi jsHjw+c4ELQFI/TkHBvvp5NaCccQoNwwW9lnIeOik
Db2lwWg56G6LiYfpVBss5+0OW2jXlq9CdNw1KLSdq+kMtZjdVr8+iQ9gsqxvskCAwEAAaAAMA0GC
SqGMIIBpjCCAQ8CAQAwZjELMAkGA1UEBhMCREUxEDA0BgNVBAGTB0dlcm1hbnkxETAPBgNVBACTC
EVyYzCBnzANBkgqhkiG9w0BAQEFAA0BjQAwgYkCgYEAz+hQRsqDWRLvmV4YD5+JaQEXn5qqJeyzk
fg2PQoU2VPgHID0VnyTPt8r/t4uFk8p1NxjYkC4
-----END NEW CERTIFICATE REQUEST-----
```

Now submit this CSR to the desired CA. For example, VeriSign offers an online order form that contains a field into which the text from the CSR can be simply pasted. Other CAs may do this differently – please ask the CA for help if there is any question, or if anything is unclear. After the CA has received the CSR, it will first verify that the requester or company is indeed genuine, that is if the content of the certificate can or cannot be trusted. This usually happens using methods such as making phone calls, checking company registrations, or other types of research, and may take a few days. Once the CA has verified the validity of the certificate, it will either be returned as a signed certificate, or instructions on how to obtain the signed certificate will be supplied.

18.3.4 Installing the Signed Server Certificate

The signed certificate received back from the CA must be in X.509 format; either in binary or Base64 encoded form (please contact the CA if the certificate received does not match either of these formats). Once the certificate has been received, store it into a file (usually “*.cer”). Then execute the following command:

```
[maestro_install_folder]/java/bin/keytool -import
-alias NAME -file INFILE -keystore KEYFILE -trustcacerts
```

with the following replacements:

NAME: The name of the certificate. This must be the name of the certificate that the CSR was made for (the same name that was specified when the certificate and the CSR were created).

INFILE: The file that contains the reply from the CA with the signed certificate.

KEYFILE: The keystore file in which the certificate is stored. This can either be a relative or a full path name.

The command will load the certificate from the given file, check the signature of the signer (the CA) against a trusted root certificate of the CA and, if the signature matches, replace the unsigned version of the certificate that was in the keystore with the signed version.

There is one critical moment here – when the tool tries to check the signature against a trusted root certificate of the CA: If this check cannot be made because such a trusted root certificate of the CA cannot be found, the tool will abort with an error message. In this case, obtain a trusted root certificate from the CA first (see Section 18.3.5 [Installing a Trusted Root Certificate](#) for more details) and then repeat the import step described above. Java is already shipped with trusted root certificates of certain CAs, like VeriSign and Thawte. For other CAs, obtain and install a root certificate first.

18.3.5 Installing a Trusted Root Certificate

This step is only required if the signed server certificate was obtained from a CA for which a trusted root certificate is not already shipped with Java. An error message during the import of the signed certificate will occur if this is the case. The required root certificate should be available from the CA. The certificate must be stored in a file, either in “DER encoded binary X.509” or “Base-64 encoded X.509” format. If there is access to such a certificate file, import it into the keystore with the trusted root certificates by executing the following command:

```
[maestro_install_folder]/java/bin/keytool -import  
-alias NAME -file INFILE -keystore  
[maestro_install_folder]/java/lib/security/cacerts
```

with the following replacements:

NAME: The name to be given to the certificate in the keystore. This name is not important for anything, except for recognition later. In addition, this name must not yet be in use in the keystore.

INFILE: The file in which the X.509 certificate from the CA is stored.

The password of the default keystore file will be queried for, which should have been set to something other than its default “changeit” earlier. See Section 18.3.1 [Securing the Trusted Root Certificate Keystore](#) for more information. The command will present the details of the certificate to be imported in a way similar to Figure 18-2.

Figure 18-2 Imported Certificate

```

Owner: OU=For VeriSign authorized testing only. No assurances (C)VS1997,
OU=www.verisign.com/repository/TestCPS Incorpor. By Ref. Liab. LTD., O="VeriSign, Inc"
Issuer: OU=For VeriSign authorized testing only. No assurances (C)VS1997,
OU=www.verisign.com/repository/TestCPS Incorpor. By Ref. Liab. LTD., O="VeriSign, Inc"
Serial number: 52a9f424da674c9daf4f537852abef6e
Valid from: Sun Jun 07 02:00:00 GMT+02:00 1998 until: Wed Jun 07 01:59:59
GMT+02:00 2006
Certificate fingerprints:
    MD5:  40:06:53:11:FD:B3:3E:88:0A:6F:7D:D1:4E:22:91:87
    SHA1: 93:71:C9:EE:57:09:92:5D:0A:8E:FA:02:0B:E2:F5:E6:98:6C:60:DE
Trust this certificate? [no]:

```

The presentation contains details about the certificate, but these could have been forged. It also contains the certificate's fingerprints, which can be used to verify that the certificate has not been falsified. For example, if the certificate was emailed (thus giving a potential attacker the possibility to "catch" the email before it reaches its destination, and replacing the certificate therein with his own certificate for a future "man-in-the-middle" attack), then it is advisable to call the responsible person from the CA, to verify the fingerprint of the certificate over the phone.

Once the certificate is believed to be genuine, answer **[Yes]** and RETURN to the question "Trust this certificate?" After this is done, the certificate is installed as a new trusted root certificate in the default Java root certificate store, and can now be used to import server certificates signed by the CA from which the root certificate was obtained, as described in Section 18.3.4 [Installing the Signed Server Certificate](#).



Tip: Internet Explorer comes (as many browsers do) with an extensive list of trusted root certificates. It also allows those certificates to be exported to a file in the X.509 format required for import by "keytool". Therefore, if a CA is chosen to sign the server certificate for which there is no trusted root certificate already in the Java default keystore, it is very easy to locate a root certificate by exporting it from Internet Explorer as described below. This description is for IE 5.0, 5.5 and 6.0; other versions may vary.

Go to Tools > Internet Options... > Content tab > Certificates... > Trusted Root Certification Authorities tab, and look for a matching root certificate (many CAs have several of these). This might have to be done by trial-and-error until a matching certificate is found. Select the certificate and click **[Export]**. In the Export wizard, choose either **DER encoded binary X.509 (.CER)** or **Base-64 encoded X.509 (.CER)** and supply a suitable filename. Next, complete the export. The file that is exported can then be imported into Java's default keystore as described above.

18.3.6 Making LISTSERV Maestro Aware of the Server Certificate

Once you have imported the signed server certificate into your `keystore` file, you now need to make the LISTSERV Maestro server aware of this certificate, as a last step of securing your server.

On the server that you want to secure with SSL, edit the file “tomcat.ini”:

```
[maestro_install_folder]/conf/tomcat.ini
```

To enable SSL, you need to add a `SecureServer` entry with the value of “true”, like this:

```
SecureServer=true
```



Note: Setting this entry to “false” or not specifying the entry at all, will have the effect that SSL will be disabled, in which case the description below does not apply either.

In addition to enabling SSL as described above, you also need to specify the information about the keystore file in which the server certificate is stored. You do so with the following two entries:

```
KeystoreFile=KEYSTORE_PATH
```

```
KeystorePassword=PASSWORD
```

with the following replacements:

- **KEYSTORE_PATH:** Replace with the absolute path to the keystore file (on Windows including drive letter) in which the signed certificate can be found. You can not use a relative path name but must supply the full path to the file.

You can store the keystore file itself in any place that seems appropriate, but the [maestro_install_folder]/conf folder seems like a good choice.

- **PASSWORD:** Replace with the password that you used for the keystore (as explained earlier, you must have used the same password for the certificate itself too).



Security Issue: As you see, the password to the keystore and the certificate therein is included as plain text in this file. This can be a security breach, if unauthorized persons have access to this file. You should therefore employ the appropriate operating system or file system security measures, so that only authorized persons can access the `tomcat.ini` file (since it is a file integral to the functioning of the server, you should have done this anyway, because tampering with this file, or other files in the installation folder, may stop LISTSERV Maestro from working properly).



You may also have to edit the `SecurePort` entry in the same `tomcat.ini` file. This entry specifies the port to be used for receiving client requests when using HTTPS protocol.

If this entry is not specified, then the default HTTPS port 443 will be used.

However, if the `SecurePort` entry is specified, it overrides this default behavior.

Therefore, once you have enabled SSL, check if the port entry matches the now enabled HTTPS protocol: If a `SecurePort` entry is present, does it specify the correct port you want to use? If not, either remove the entry (to use the default port 443) or change its value to your desired port.

If you do change the `SecurePort` entry, please see also Section 14.2.2 [Configuring the HTTPS Port](#) about other issues that apply when changing the HTTPS port (for example, adjusting certain HUB settings).

LISTSERV Maestro is now prepared for SSL access. Start (or restart) LISTSERV Maestro to make the changes effective.

You can access LISTSERV Maestro normally, only now you need to use “**https:**” URLs instead of the standard “http:”, because once you have enabled and configured SSL access as described above, the communication with **all** components on this server may now only happen via HTTPS anymore. Normal access to this server via HTTP is no longer possible.



Note: The above procedure describes what you need to do if you want to secure the whole server with SSL, i.e. all components that are installed on that server. If instead you only want to secure some components, then see Section 18.3.8 [Mixing SSL and Non-SSL Access on One Server \(Mixed Mode\)](#) for details of how to configure such a mixed mode.

18.3.7 Securing a Server with Multiple Host Names

Under normal circumstances, SSL does not work if the server has several different host names, which are mapped to the same IP-address. This is a technical limitation of SSL, not a specific limitation of LISTSERV Maestro.

The reason for this is, that SSL is supposed to encrypt all data that is being transferred between the server and the browser. Therefore, the SSL handshake (through which server and browser negotiate the encryption details) has to happen as the very first thing once a SSL connection has been established, before any other data is transferred (otherwise that data would not be encrypted). This means that not even the initial request itself (which contains the hostname that the request is directed to) has been transferred at this point.

Therefore, at the moment the SSL handshake is happening, the server does not yet know which host name the browser is trying to access. The server only knows the IP-address that the browser is connecting to. Because of this, the server must determine the certificate that is to be used for the SSL encryption based only on the IP-address. It can not select the certificate based on the host name that is contained in the request.

But if the IP-address is mapped to several host names, then the server will likely select the wrong certificate. Because if there are several certificates for the same IP-address (as is the case if there are several certificates for different host names, which however all are mapped to the same IP-address), then the server will simply pick one of them, with no guarantee that this will be the certificate that matches the host name that is contained in the request.

Therefore, it is usually not possible to SSL-secure a server that has several different host names, which are all mapped to the same IP-address.

If you have such a server with multiple host names, you need to assign several IP-addresses to the server so that each host name can have its own dedicated IP-address, i.e. each IP-address that is assigned to the server must be used by only one of the server's host names.

Then, you can create and install certificates for all these host names, and the individual IP-addresses will make sure, that for each SSL request that is directed to one of the host names, the correct certificate that matches the IP-address of that host name will be selected.



How you do this depends on your scenario:

If you have multiple host names for the server, but want to SSL-secure the server for all these host names (and for all Maestro components), then you do this in the same fashion as described in the previous sub-sections, only that you need to have a separate IP-address for each host name, and you need to obtain a separate server certificate for each host name too, where you store all these server certificates in the same keystore file.

But if you want to SSL-secure only some of the host names and leave some unsecured, then you need to run the server in mixed mode, where you SSL-secure access to some of these IP-addresses, but not to others (see Section 18.3.8 [Mixing SSL and Non-SSL Access on One Server \(Mixed Mode\)](#) for details).



18.3.8 Mixing SSL and Non-SSL Access on One Server (Mixed Mode)

The requirement for running the LISTSERV Maestro server in mixed mode (partially SSL-secured and partially not) can arise from several scenarios:

- If there are several components installed on the server (for example, all three) but only some of them are supposed to be SSL-secured (for example, LUI/HUB) but others are not (for example, TRK).
- If the server has several host names, and only some of them will be SSL-secured but others will not.
- Any mix of the two above scenarios. For example: A server with several host names, some of which will be SSL-secured and some not, and even for the SSL-secured host names, only some of the installed LISTSERV Maestro components will be secured but others not.

There are a few pre-conditions for configuring SSL/non-SSL mixed mode:

- If you want to SSL-secure some host names, but not others, then for each host name that you want to secure, you need a separate IP-address. This means that this IP-address must only be used by this one host name. A secured host name is *not* allowed to share an IP-address with another secured host name, and it is *not* allowed to share an IP-address with one of the unsecured host names (but the unsecured host names may share one IP-address among them).

For example: Assume that you have five host names, of which two will be secured and three not. For this, you need a minimum of three IP-addresses: One to be used solely by the first secured host name, a second IP-address to be used solely by the second secured host name, and a third IP-address to be shared by the three unsecured host names.

- If you want to secure some of the installed LISTSERV Maestro components but not others, then you need at least two host names with unique IP-addresses: One of these host name/IP-address pairs would be used for the components that are to be secured, and the other pair used for the components that are not be secured.

For example: Assume that you have LUI, HUB, and TRK on the same server and want to secure LUI and HUB, but not TRK. You will need two host names with two

unique IP-addresses: One pair to be used by LUI/HUB (secured) and one pair to be used by TRK (unsecured).

Depending on your scenario and according to the above pre-conditions you will have a list of host names that you want to secure, each of which must have a unique IP-address that is not used by any other host name. In the description below, we will call these the “secured host names” and “secured IP-addresses”.

With these secured host names and secured IP-addresses, you can then configure the SSL/non-SSL mixed mode more or less in the same fashion as when securing the whole server (as described in the previous sub-sections), except for the following differences:

- In the `tomcat.ini`, instead of supplying “`SecureServer=true`” (as you would if you wanted to secure the whole server), you now supply a comma-separated list of all the secured IP-addresses, like this:

```
SecureServer=SECURED_IP_ADDRESS1,SECURED_IP_ADDRESS2,...
```

For example:

```
SecureServer=192.168.0.15,192.168.0.16,192.168.0.43
```



Note: You must use this method of specifying the secured IP-addresses even if you have only a single secured IP-address.

- Also, in the `tomcat.ini`, check the values of the settings `Port` (which defines the HTTP port used by all unsecured host names) and `SecurePort` (which defines the HTTPS port used by all secured host names). If any of these settings is not present, then the default port for the corresponding protocol will be used, so if you want to use a different port, make sure to specify the corresponding setting.
- Obtain a server certificate for each of the secured host names, but do not store these server certificates in one and the same keystore file. Instead, you need to have a separate keystore file for each certificate. This means that you need one separate keystore file for each secured host name containing the certificate for that host name.

These keystore files must all observe the following constraints:

- All keystore files must use the same password. This is also the password that you need to specify in the `tomcat.ini` as the value of the `KeystorePassword` setting.
- All keystore files must use the same keystore type (as specified by the advanced `tomcat.ini` setting `KeystoreType`).
- All keystore files must be stored in the same folder.
- All keystore files must be named according to the following naming pattern:

```
IP_ADDRESS-COMMON_FILENAME
```

where `IP_ADDRESS` is the secured IP-address that corresponds to the secured host name to which the keystore file belongs, and `COMMON_FILENAME` is a freely selectable file name, which however must be the same for all keystore files.

For example, let's assume we have three keystore files for secured host names with the secured addresses "192.168.0.15", "192.168.0.16" and "192.168.0.43". And for the common file name we chose "certificate.keystore". Then we would have to create three keystore files with the following full file names:

```
192.168.0.15-certificate.keystore
```

```
192.168.0.16-certificate.keystore
```

```
192.168.0.43-certificate.keystore
```

- For the `KeystoreFile` setting in the `tomcat.ini`, supply the following value:

```
KeystoreFile=PATH_TO_FOLDER_OF_KEYSTORE_FILES/COMMON_FILENAME
```

So for our example of above, assuming that the keystore files are stored in a folder called "C:/Maestro/conf", the entry would look like this:

```
KeystoreFile=C:/Maestro/conf/certificate.keystore
```

Note, that you specify only a single "KeystoreFile" entry in the `tomcat.ini`, even though there are several keystore files! When specifying this entry, you supply only the path and the common filename. LISTSERV Maestro will then automatically generate the actual keystore filenames from this setting by inserting each secured IP-address (with the separating "-" character) between the path and the common file name part. So for each secured IP-address, the actual filename will become something like:

```
C:/Maestro/conf/192.168.0.15-certificate.keystore
```

which then matches the file names that we gave our keystore files (see above).

- Finally, check the **URL Settings** in the HUB (on default level and/or group or user level) to check that the protocol, hostname, and port that is specified there matches your SSL/non-SSL configuration. Access URLs to secured host names must use "https://" while access URLs to unsecured host names must use "http://".

18.3.9 Supported SSL Cipher Suites

The following table lists the SSL ciphers that are supported by LISTSERV Maestro. It also shows which ciphers are enabled as allowed ciphers by default, if the `SSLCiphers` setting in the `tomcat.ini` is not defined. If you want to define a different set of allowed SSL ciphers, add the `SSLCiphers` setting to the `tomcat.ini` and specify the cipher-names in a comma separated list as the value of the setting.

Table 18-1 Supported SSL Ciphers

Name	Allowed by Default
SSL_RSA_WITH_RC4_128_MD5	Yes
SSL_RSA_WITH_RC4_128_SHA	Yes
TLS_RSA_WITH_AES_128_CBC_SHA	Yes
TLS_DHE_RSA_WITH_AES_128_CBC_SHA	Yes

Name	Allowed by Default
TLS_DHE_DSS_WITH_AES_128_CBC_SHA	Yes
SSL_RSA_WITH_3DES_EDE_CBC_SHA	Yes
SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA	Yes
SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA	Yes
SSL_RSA_WITH_DES_CBC_SHA	Yes
SSL_DHE_RSA_WITH_DES_CBC_SHA	Yes
SSL_DHE_DSS_WITH_DES_CBC_SHA	Yes
SSL_RSA_EXPORT_WITH_RC4_40_MD5	Yes
SSL_RSA_EXPORT_WITH_DES40_CBC_SHA	Yes
SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA	Yes
SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA	Yes
SSL_RSA_WITH_NULL_MD5	
SSL_RSA_WITH_NULL_SHA	
SSL_DH_anon_WITH_RC4_128_MD5	
TLS_DH_anon_WITH_AES_128_CBC_SHA	
SSL_DH_anon_WITH_3DES_EDE_CBC_SHA	
SSL_DH_anon_WITH_DES_CBC_SHA	
SSL_DH_anon_EXPORT_WITH_RC4_40_MD5	
SSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA	
TLS_KRB5_WITH_RC4_128_SHA	
TLS_KRB5_WITH_RC4_128_MD5	
TLS_KRB5_WITH_3DES_EDE_CBC_SHA	
TLS_KRB5_WITH_3DES_EDE_CBC_MD5	
TLS_KRB5_WITH_DES_CBC_SHA	
TLS_KRB5_WITH_DES_CBC_MD5	
TLS_KRB5_EXPORT_WITH_RC4_40_SHA	
TLS_KRB5_EXPORT_WITH_RC4_40_MD5	
TLS_KRB5_EXPORT_WITH_DES_CBC_40_SHA	
TLS_KRB5_EXPORT_WITH_DES_CBC_40_MD5	

Section 19 Tracking and Recipient Profiles

Among the four tracking types, LISTSERV Maestro offers two types that involve recipient profiles – personal tracking and anonymous tracking. With personal tracking, each recipient is identified uniquely by a recipient ID that can be traced back to the data associated with this recipient, (the recipient's profile). With anonymous tracking, each recipient is identified with an anonymous ID that cannot be traced back to the actual recipient data, but only to an anonymous profile. This is usually a subset of the actual recipient data that contains only anonymous data, no personal data such as name or address is included.

When anonymous tracking is chosen, LISTSERV Maestro always creates and stores an anonymous profile for each recipient. For higher efficiency, if several recipients have the same anonymous profile, only one profile entry is created and this is shared by all of the recipients. The anonymous ID is then included in the tracking data and maps to one of these anonymous profiles stored in LISTSERV Maestro.

The storage of personal profiles is very similar. For each recipient, a profile entry with this recipient's data is created. Usually there will be one entry for each recipient, but should several recipients happen to have exactly the same profile, only one profile entry will be generated and this will be shared by those recipients. Both anonymous and personal profiles are stored in the Maestro System database. See Section 10 [The System Database](#) for additional information.

Anonymous profiles always need to be created and stored by LISTSERV Maestro, because they simply do not exist anywhere else. However, with personal profiles, this is usually different. The personal profile of a recipient contains the full set of data associated with that recipient. It maps to one row in the uploaded recipients file (in CSV format), or to one row in the result set that was selected from the database. Each column in the row constitutes one field of the profile data, where the column headers from the uploaded file or the database table are the labels of these fields.

For personal tracking, the recipient data must also contain one column with a unique recipient ID – a column with values that can be used to identify the recipient from all other recipients.

More often than not, the recipient data already comes from some type of database. Either it was exported from the database and then uploaded as a recipients file, or the Maestro User Interface selected it directly from the database (possibly by using a database backed target group). In both cases, there is already a table in a database that contains the full recipient profiles, including the unique recipient IDs. In some cases, when the Maestro User Interface is used with an external system database, and that database happens to be the same database as the one where the recipients originally came from, the original recipient profiles exist in the same database where the Maestro User Interface will store them. Under certain circumstances, therefore, it seems redundant to allow the Maestro User Interface store personal profile information in its database, when the same information already exists in another database (or even in the same database, if the database is shared as explained above).

To avoid this circumstance, the Maestro User Interface offers an option to switch off the storing of personal profiles in the Maestro System Database. To do this, edit the following file:

```
[maestro_install_folder]/lui/lui.ini
```

Add this entry:

```
CreatePersonalProfileTables=false
```

If the entry is set to `false`, then the Maestro User Interface will not write personal profiles into its system database. If it is set to `true` (or missing, which is the *default* after installation) the Maestro User Interface will create personal profiles. Restart the Maestro User Interface after the change to make the entry effective.

The actual difference between permitting and not permitting the Maestro User Interface to create personal profiles is that if the Maestro User Interface creates personal profiles, then the match between the recipient ID that is collected with the tracking event and the corresponding recipient (that recipient's profile) can be made directly in the Maestro User Interface.

If the report type Details Report is run, the resulting table will have one entry for each recipient for which one of the events selected was registered. Optionally, with a count that details the amount of these events that were registered. One row per recipient is generated, including the recipient's profile as values in the row.

Figure 19-1 Example of Recipients Profile Data Table

```
"Count " , "ID" , "Name " , "EMail" , "Age " , "ZIP "
"5" , "fred1" , "Fred" , "fred@flintstone.com" , "52" , "12345"
"2" , "wilma1" , "Wilma" , "wilma@flintstone.com" , "45" , "12345"
etc...
```

The **Count** column is optional.

With this table, it is immediately apparent which recipients reacted to the message (and how often, if the **Count** column is included), as the details of each recipient are included in the form of a profile.

If the version of this table without the **Count** column is chosen, the same table can also be used, without any modifications, to upload the recipients list for another job (for example to send a follow-up mail to all recipients that reacted to the previous mail). The data is already in the CSV-format that the Maestro User Interface understands, and since all recipient profiles are already in the Maestro User Interface database, the profiles will not be recreated, but instead the existing profiles are reused.

In contrast, if the Maestro User Interface does not create personal profiles, then it is necessary to make the match between the recipient IDs and the actual recipients behind them with a tool outside of LISTSERV Maestro, because the Maestro User Interface does not contain the information to do so itself. To help make this match, the Maestro User Interface will output a table with the recipient IDs in question when the Details Report is run. The result is one row per recipient, with the recipient's ID as the value in the row.

Figure 19-2 Example of Recipients ID in Data Table

```
"Count" , "ID"  
"5" , "fred1"  
"2" , "wilma1"  
etc...
```

Again, the column **Count** is optional.

Here, only the ID's of the recipients that reacted (and how often, if the **Count** column was included) are apparent, but any further details regarding the recipients are not.

This data would have to be brought into context with the original source of the recipients, by whatever reporting or analysis tool preferred to discover more details about the users.

The type of handling of the personal profiles depends on the requirements of the feedback desired:

- If immediate and simple-to-get feedback is desired about the recipients who trigger the events, and there is no concern about saving storage space, (keeping redundant versions of the profiles in different databases) choose the option of permitting the Maestro User Interface to create personal profile entries. Set the INI-file entry to `true` or leave it out, which is the default after installation.
- If there is no concern about receiving feedback on the identity of the recipients quickly and there is concern about saving disk space, keeping redundant sets of data is not desired. Choose the option of not storing profile entries in the Maestro System Database by setting the INI-file entry to `false`.

The choice between allowing and not allowing the Maestro User Interface to store personal profiles in the system database is really an advanced administration feature. If there is any concern about this choice, keep the default of letting the Maestro User Interface store the profiles. Only change this setting after thoughtful consideration of the requirements and the impact this will have.

Section 20 Editing LISTSERV Maestro INI Files

The following rules apply to the `lui.ini`, `hub.ini`, `tracker.ini`, and `tomcat.ini` INI-files, which are the configuration files for the Maestro User Interface, Administration Hub, and Maestro Tracker components. (The `my.ini` configuration file of the internal MySQL database follows different rules since it's a third party product. See the MySQL documentation for details.)

- All INI-files are text files and assumed to be encoded in the default encoding for the platform being used. For most English/European installations that would be ISO-8859-1 [Latin 1 – Western European].
- In the files, every parameter occupies one line. Each line must be terminated by a line terminator (LF, CR, or CRLF). All lines in the file are processed.
- A line that contains only white space or whose first non-white space character is an ASCII “#” or “!” is regarded as a comment, and its content is ignored.
- Every line other than a blank line or a comment line describes one parameter (except if a line ends with a backslash “\”, then the following line, if it exists, is treated as a continuation line, as described below).
- A parameter always consists of a key and a value. Keys and values are separated by white space or “=” or “:”. Any white space around the separation character is also ignored.
- All remaining characters on the line become part of the associated value. Some characters which otherwise have special meanings, need to be escaped with a backslash. The ASCII escape sequences “\t”(TAB), “\n”(LF), “\r”(CR), “\\”(backslash), “\ ”(quotation mark), “\ ’”(apostrophe), “\ ”(space), and “\uxxxx” (where “xxxx” is the Unicode-value of the required character, expressed in hexadecimal format) are recognized and converted to single characters.
- In the case that the last character on a line is a “\”, then the next line is treated as a continuation of the current line; the “\” and line terminator is simply discarded, and any leading white space characters on the continuation line are also discarded and are not part of the parameter value.

Examples:

Each of the following four lines specifies the key "Truth" and the associated value "Beauty":

```
Truth = Beauty
Truth:Beauty
Truth :Beauty
Truth Beauty
```

The following three lines specify a single parameter:

```
fruits      apple, banana, pear, \
            cantaloupe, watermelon, \
            kiwi, mango
```

The key is "fruits" and the associated value is:

"apple, banana, pear, cantaloupe, watermelon, kiwi, mango"

Note that a space appears before each "\" so that a space will appear after each comma in the final result; the "\", the line terminator and leading white space on the continuation line are discarded and are not replaced by one or more other characters.

As a last example, the line:

```
cheeses
```

specifies that the key is "cheeses" and the associated value is the empty string.




20.1 Maestro User Interface INI-File Entries

The following table shows all possible entries of the `lui.ini` file for the Maestro User Interface component. For any entry that is missing in the INI-file, the corresponding default value is assumed. Changes in INI files require a restart of the component to take effect.

Table 20-1 Maestro User Interface INI-File Entries

Entry Key	Description
AddJobIdToAllMessages	Defines if the LISTSERV Maestro job ID should be included in the message ID of all outgoing emails or not. Set to either <code>true</code> or <code>false</code> . Default: <code>false</code>
AllowCharsetChoice	Defines if the user is allowed to change the content charset on a job-by-job basis, or if he has to accept the default charset (see "DefaultMailCharset"). Default: <code>true</code> See Section 27.1 Defining the Default Mail Charset .
AllowISO-i-Mails	Defines if the special bi-directional charsets "ISO-8859-6-i" and "ISO-8859-8-i" will be used in outgoing mail instead of their normal Iso-8859 counterparts. Default: <code>true</code> See Section 27.2 Allowing or Disallowing Bi-Directional Character Sets .

AllowPersistentCookies	<p>Defines if the Maestro User Interface is allowed to store persistent cookies in the client browser so that the user interface may remember certain user settings between two sessions. If set to “false”, then these settings will be stored as session-cookies only and will be forgotten when the browser is closed.</p> <p>Default: true</p>
ChangeLog	<p>Specifies whether or not a change log of subscriber activities (for hosted lists and datasets) is to be created.</p> <p>(Optionally) If a change log file is used, then this defines the time period for each log (i.e. daily, weekly, monthly, or yearly).</p> <p>See Section 13.2 Subscriber Activity Change Log.</p> <p>Default: false</p>
ClickThroughURL	<p>Path-part of the click-through tracking URL used for URLs without passing of merged parameters.</p> <p>Default: /trk/click</p>
ClickThroughPPURL	<p>Path-part of the click-through tracking URL used for URLs with passing of merged parameters.</p> <p>Default: /trk/click</p>
CreatePersonalProfileTables	<p>Specify whether full personal profiles should be stored in the system database when personal tracking is used.</p> <p>Default: true</p> <p>See Section 19 Tracking and Recipient Profiles.</p>
DashboardReportLifetimeBase	<p>Defines the basis for the calculation of the dashboard report lifetime, in minutes. For efficiency’s sake, a dashboard report is refreshed again only after its lifetime has expired. The lifetime until the next refresh is re-calculated after each refresh as a random value in the range from “baseTime” to “baseTime*1.5”.</p> <p>Note: Tracking event reports do not use this base-time, as they determine their lifetime based on the event transfer interval instead.</p> <p>Default: 4 (i.e. by default the actual dashboard lifetime is a random value in the range between 4 and 6 minutes.)</p>

DefaultCustomizationLanguage	Defines the language that is to be used for the membership area subscriber pages of all datasets that do not have their own membership area language defined. The language must be specified using its corresponding two-letter ISO language code (in lower case), as defined by the alpha-2 code of ISO-639-1 (see http://www.loc.gov/standards/iso639-2/php/English_list.php). Default: "English (Default)"
DefaultMailCharset	Defines the charset that is to be used as the charset of the content for newly created jobs that are not copies of existing jobs. May or may not be changed by user (see "AllowCharsetChoice"). Default: ISO-8859-1 See Section 27.1 Defining the Default Mail Charset .
DistributeChunkSize	Defines if distribute chunking will be performed and with which chunk size. A chunk size of zero or less means that no distribute chunking will be performed. A positive non-zero chunk size means that chunking will be performed with chunks of the given size, but with a minimum size of 10000 (i.e. any positive chunk size less than 10000 will be treated as if 10000 had been specified). Default: 100000
 DistributeBaseTimeout	Defines the default timeout for socket connections to LISTSERV, in minutes. Any value less than 1 will be treated as 1, i.e. 1min is the minimum possible timeout. Default: 30
ExpiredConfirmCleanupInterval	Determines how often the system looks for (and cleans out) expired unconfirmed subscriptions to hosted tables and lists (i.e., subscriptions that were not confirmed during the given expiration interval – see INI-file parameter <code>OptInConfirmationExpiration</code>). Specified as a number, which defines the cleanup interval length in "hours" (i.e., every "N" hours the system checks for expired subscriptions and/or external profile changes and cleans them out). Default: 1
 ExternalHostName	<i>This INI parameter is no longer supported; it has been replaced by a setting in the HUB.</i>
 ExternalHTTPPort	<i>This INI parameter is no longer supported; it has been replaced by a setting in the HUB.</i>

New	ExternalProfileEditConfirmExpiration	<p>The expiration time of unconfirmed external changes to membership area and list subscription profiles, i.e. the maximum amount of time after the user has supplied the changed values on the external profile edit page, during which the user must confirm the changes, otherwise the changes expire (are removed from the temporal storage area). Specified as a number, which defines the expiration time in "hours".</p> <p>Default: 96</p>
New	ExternalProtocolIsHTTPS	<p><i>This INI parameter is no longer supported; it has been replaced by a setting in the HUB.</i></p>
	Home	<p>Home folder in which work-files are kept.</p> <p>Default: subfolder "lui" in installation folder</p>
	HostName	<p>The host name of the local machine that is to be used whenever a name is required to identify the local host (and no other INI-parameter overrides this for a specific purpose).</p> <p>Default: normal host name (canonical host name) of the server running the Maestro User Interface.</p>
	HubContext	<p>Context-path part of the user interface access URLs for the Administration Hub component.</p> <p>Default: hub</p>
New	HubExternalHostName	<p><i>This INI parameter is no longer supported; it has been replaced by a setting in the HUB.</i></p>
New	HubExternalHTTPPort	<p><i>This INI parameter is no longer supported; it has been replaced by a setting in the HUB.</i></p>
New	HubExternalProtocolIsHTTPS	<p><i>This INI parameter is no longer supported; it has been replaced by a setting in the HUB.</i></p>
	HubRMIPort	<p>Internal communication port (RMI-Port) of the Administration Hub server.</p> <p>Default: 1099</p> <p>See Section 14.2.3 Configuring the Internal Communication Port.</p>
	LogFolder	<p>Defines the folder under which the Maestro User Interface component stores the log files.</p> <p>Default: The "logs" subfolder of the LUI home folder (i.e., by default, the subfolder "lui/logs" of the installation folder).</p>
	MaintenanceMode	<p>Defines if the Maestro User Interface component will run in maintenance mode or not.</p> <p>Default: false</p> <p>See Section 6.3 User Restrictions.</p>

OpenUpURL	Path-part of the open-up tracking URL. Default: /trk/open
OptInConfirmExpiration	The expiration time for unconfirmed subscriptions to hosted datasets and lists – i.e., the amount of time after the initial subscription request during which the subscriber may confirm the subscription (“double opt-in”) before the subscription request expires and is removed. Specified as a number, which defines the expiration time in hours. Default: 48 See the LISTSERV Maestro Data Administrator’s Manual for more information about subscription confirmations.
RegistryHubHost	Host name of the server with the Administration Hub component. Default: localhost
RemoteAdminPassword	Password for remote log file access. Default: <i>none</i> (no remote log file access allowed)
RestoreBackup	Path name of the folder containing the backup that shall be restored during the next startup. Note: This key will be automatically removed from the INI-file during the next startup. Default: <i>none</i> See Section 11.7 Restoring a Backup .
RMIPort	Internal communication port (RMI-Port) of the Maestro User Interface server. Default: 1099 See Section 14.2.3 Configuring the Internal Communication Port .
ShowEventsCountOnDashboard	Define if the “Currently in the system” section of the dashboard is supposed to display the current count of events in the system or not. On systems with a lot of events, the calculation of this count may be slow, thus slowing down the dashboard display. In this case it may be a good idea to disable the display of this count by supplying “false”. Default: true
TrackerHost	Host name of the server with the Maestro Tracker component. Default: localhost See Section 24.2.3 Moving the Maestro Tracker Component to Another Server .

TrackerRMIPort	Internal communications port (RMI-Port) of the Maestro Tracker server Default: 1099 See Section 14.2.3 Configuring the Internal Communication Port .
TrkExternalProtocolIsHTTPS	<i>This INI parameter is no longer supported; it has been replaced by a setting in the HUB.</i>
UnsubscribedCleanupInterval	Determines how often the system looks for (and cleans out) subscribers of hosted LISTSERV lists that have unsubscribed by email (i.e. not by the LISTSERV Maestro web interface). Specified as a number, which defines the cleanup interval length in “minutes” (i.e., every “N” minutes the system checks for unsubscribed subscribers and cleans out their data.) Default: 10 See the LISTSERV Maestro Data Administrator’s Manual for more information about subscription confirmations.



20.2 Administration Hub INI-File Entries

The following table shows all possible entries of the `hub.ini` file for the Administration Hub component. For any entry that is missing in the INI-file, the corresponding default value is assumed. Changes in INI files require a restart of the component to take effect.

Table 20-2 Administration Hub INI-File Entries

Entry Key	Description
Home	Home folder in which work-files are kept. Default: subfolder “hub” in installation folder
HostName	The host name of the local machine that is to be used whenever a name is required to identify the local host (and no other INI-parameter overrides this for a specific purpose). Default: normal host name (canonical host name) of the server running the Administration Hub.
LiteMode	Defines if LISTSERV Maestro will run in lite-mode or not. Only in lite-mode will LISTSERV Maestro accept lite-LAKs during login to LUI. Set to <code>true</code> for lite-mode or <code>false</code> for full-mode. Default: <code>false</code>

LogFolder	Defines the folder under which the Administration Hub component stores the log files. Default: The "logs" subfolder of the HUB home folder (i.e., by default, the subfolder "hub/logs" of the installation folder).
RegistryDomain	The domain name with which the Administration Hub component stores its settings in its own registry. Default: HUB
RemoteAdminPassword	Password for remote log file access. Default: <i>none</i> (no remote log file access allowed) See Section 13.1 Remote Log Access .
RMIPort	Internal communication port (RMI-Port) of the Administration Hub server. Default: 1099 See Section 14.2.3 Configuring the Internal Communication Port .

20.3 Maestro Tracker INI-File Entries

The following table shows all possible entries of the `tracker.ini` file for the Maestro Tracker component. For any entry that is missing in the INI-file, the corresponding default value is assumed. Changes in INI files require a restart of the component to take effect.

Table 20-3 Maestro Tracker INI-File Entries

Entry Key	Description
FixUserAgentTable	Set this key to "true" if you want Tracker to check the user-agents table for consistency and to repair it (if there is any file corruption found). Tracker will notice the key during the next startup and perform the check/repair during the startup. This key is automatically removed from the INI-file during startup. Default: none Use this only if instructed to by L-Soft Support.
Home	Home folder in which work-files are kept. Default: subfolder "trk" in installation folder
HostName	The host name of the local machine that is to be used whenever a name is required to identify the local host (and no other INI-parameter overrides this for a specific purpose). Default: normal host name (canonical host name) of the server running Maestro Tracker.
HubRMIPort	Internal communication port (RMI-Port) of the Administration Hub server. Default: 1099 See Section 14.2.3 Configuring the Internal Communication Port .

LogFolder	Defines the folder under which the Maestro Tracker component stores the log files. Default: The "logs" subfolder of the TRK home folder (i.e., by default, the subfolder "trk/logs" of the installation folder).
RegistryDomain	The domain name with which the Maestro Tracker component stores its settings in the Administration Hub registry. Default: TRK
RegistryHubHost	Host name of the server with the Administration Hub component. Default: localhost
RemoteAdminPassword	Password for remote log file access. Default: <i>none</i> (no remote log file access allowed) See Section 13.1 Remote Log Access .
RMIPort	Internal communication port (RMI-Port) of the Maestro Tracker server. Default: 1099 See Section 14.2.3 Configuring the Internal Communication Port .


20.4 Tomcat INI-File Entries





The following tables show all possible entries of the `tomcat.ini` file for Tomcat. For any entry that is missing in the INI-file, the corresponding default value is assumed. Changes in INI files require a restart of the component to take effect.

20.4.1 Basic Tomcat Configuration Parameters

The following table shows the basic entries of the `tomcat.ini` file for Tomcat.

Table 20-4 Basic Configuration for Tomcat INI-File Entries

Entry Key	Description
AdditionalHost. <i>N</i>	Defines an additional host to be used for added custom content. See Section 26.2 Adding Custom Content to the Tomcat Server for details. Default: By default there are no additional hosts.
BindAddress 	Defines the local IP-addresses and ports that the server shall bind to, i.e. the server will accept connections only if directed to one of these IP addresses, in combination with the correct port. Default: By default the server binds to all addresses of the local computer (i.e. accepts connections on all these addresses), on the default port that is specified by the Port setting (see below). See Section 15 Defining the IP Addresses and HTTP Ports .

MaxPostSize		The maximum size (in bytes) that the server will accept in a POST request (i.e. the maximum size of bytes that be uploaded to the server). The limit can be disabled by setting to a value less than or equal to 0. Default: 1 (i.e. the limit is disabled by default).
Port		The HTTP port used to accept client connections. Default: 80.
Restrict.CONTEXT.N		Restrict access to component specified by "CONTEXT" (HUB, LUI, TRK, and LIST) to certain IP addresses. See Section 17.1 IP Address Restrictions .
ShutdownPort		The port used listen for shutdown requests. Default: 8007
SSL-Related Entries (See Section 18 Securing Access with SSL)		
SecurePort		Defines the HTTPS port used to accept client connections. Default: 443
SecureServer		Defines if the server is supposed to be secure (<code>true</code>) or not (<code>false</code>), or if there shall be a mixed mode (some IP addresses are secure, some are not). A secure server uses the HTTPS protocol; a non-secure server used the normal HTTP protocol. Default: <code>false</code>
KeystoreFile	 	Only used if "SecureServer=true" or if "SecureServer" defines a mixed mode. The name of the keystore-file that contains the server certificate to be used for the HTTPS protocol. The file must exist and must contain a usable server certificate. The given filename must include the full absolute path to the file (for Windows, include the drive letter). In mixed mode, the filename is acutally only a base name and the actual filenames for several keystore files will be generated from this base name. Default: <code>[install_folder]/conf/keys.keystore</code>
KeystorePassword		Only used if "SecureServer=true" or if "SecureServer" defines a mixed mode. The password used to access the keystore-file(s) as defined by the "KeystoreFile" entry. Default: <code>changeit</code>

20.4.2 Advanced Tomcat Configuration Parameters

The following table shows the advanced entries of the `tomcat.ini` file for Tomcat. Do not change these unless instructed by L-Soft Support.

Table 20-5 Advanced Configuration for Tomcat INI-File Entries

Entry Key	Description
AcceptCount	The maximum queue length for incoming connection requests when all possible request processing threads are in use. Any requests received when the queue is full will be refused. Default: 100
AJPConnector	Adds a special Tomcat connector for the AJP protocol, which is necessary to access LISTSERV Maestro through an IIS instance. See Section 32 Sharing a Server with IIS . Default: do not configure an AJP connector.
AllowTrace	Specifies whether the TRACE HTTP method is enabled (<code>true</code>) or not (<code>false</code>). Default: <code>false</code>
BufferSize	The size (in bytes) of the buffer to be provided for input streams servicing incoming requests. Default: 2048
ConnectionLinger	The number of milliseconds during which the sockets used by the server will linger when they are closed. Socket linger can be disabled by setting this entry to "1". Default: 1 (disabled by default)
ConnectionTimeout	Defines the number of milliseconds the server will wait, after accepting a connection, for the request URI line to be presented. Default: 60000 (60 seconds)
DefaultContext	Defines the default context. See Section 26 Adding Content to the Tomcat Server for details. Default: none
MaxKeepAliveRequests	The maximum number of HTTP requests which can be pipelined until the connection is closed by the server. Setting this entry to "1" will disable HTTP/1.0 keep-alive as well as HTTP/1.1 keep-alive and pipelining. Setting this entry to "-1" will allow an unlimited amount of pipelined or keep-alive HTTP requests. Default: 100
MaxSpareThreads	The maximum number of unused request processing threads that will be allowed to exist until the thread pool starts stopping the unnecessary threads. Default: 50
MaxThreads	The maximum number of request processing threads to be created by this server (i.e. the maximum number of simultaneous requests that can be handled). Default: 200

New

MinSpareThreads	<p>The number of request processing threads that will be created when the server is first started. The server will also make sure that it has the specified number of idle processing threads available. This entry should be set to a value smaller than that for the “MaxThreads” entry.</p> <p>Default: 10</p>
SSL-related Entries (See Section 18 Securing Access with SSL)	
KeystoreType	<p>Only used when “SecureServer=true”.</p> <p>Defines the type of the keystore-file. Possible types are “JKS” or “PKCS12”. Usually you should use keystore files of the “JKS” type.</p> <p>Default: JKS</p>
SSLCiphers	<p>Only used when “SecureServer=true”.</p> <p>Defines the SSL encryption ciphers that LISTSERV Maestro is allowed to use. As the value, specify a comma separated list of the allowed cipher names as shown in Section 18.3.9 Supported SSL Cipher Suites. Make sure to specify ciphers that can be used together with the configured SSLProtocol (see below).</p> <p>Default: See Section 18.3.9 Supported SSL Cipher Suites</p>
SSLProtocol	<p>Only used when “SecureServer=true”.</p> <p>Defines the version of the SSL protocol to use. Possible versions are “TLS” and “SSL”. It is recommended to always use the “TLS” protocol.</p> <p>Default: TLS</p>


 New

Section 21 Authenticating Message Origin with DomainKeys Signatures

LISTSERV Maestro lets you use DomainKeys signatures to authenticate that the messages (sent for a specific email job) do indeed originate from the domain in the “From:” address. Major ISPs already check every incoming mail to see if it is signed with a valid DomainKeys signature. Once DomainKeys has become an accepted standard for message origin verification, the current policy of only informing the recipient about the DomainKeys verification result in an additional header entry may change, and an ISP may opt to not even deliver the message to the recipient or to mark it as coming from an unsure origin. Therefore, in order to achieve good deliverability, signing messages with a valid DomainKeys signature can become more important in the future.

Support for DomainKeys signatures in LISTSERV Maestro works on three levels:

- LISTSERV as the mail distribution engine must be configured to support DomainKeys signatures for certain address domains, which requires creating a valid private/public RSA key pair and additional configuration of LISTSERV parameters. See the LISTSERV documentation for further details about setting up DomainKeys at a LISTSERV host.
- The LISTSERV Maestro Administration Hub lets you enable or disable DomainKeys signatures on the application default level and the group/single user level.
- The LISTSERV Maestro User Interface contains settings for defining the sender of a certain email job.

To define application-wide default settings for DomainKeys signatures, open the Administration Hub, click on the **Global Settings** menu, select **Maestro User Interface**, and then finally select **Default DomainKeys Settings**. The DomainKeys Settings screen opens, allowing you to define the default behavior for DomainKeys signatures.

This screen also lets you define whether or not users are allowed to change the DomainKeys signature settings for each job.

Digitally signing email messages following the DomainKeys standard is a means to assert that the message indeed originated from the domain that is claimed in the “From:” address. The digital signature is created for the whole message, which has the additional benefit that the recipient (once he or the receiving MTA has verified the signature) can be sure that the message has not been modified on its path from the sender to the recipient.

Before enabling DomainKeys support in the application, bear in mind that if DomainKeys signatures are enabled for a mail job, then all messages from the mail job must be run through a signature computation, which in most cases slows down mail job delivery.



Note: Changing the settings on this screen only applies to mail jobs that have not yet been authorized for delivery.

Figure 21-1 The DomainKeys Settings Screen

DomainKeys Settings

Define if e-mails that are sent by the user shall use DomainKeys signing by default. Additionally, define if the user is allowed to decide about DomainKeys signing (using the default value defined above).

Settings for Mail Jobs

Default Setting

Use inherited value: Yes, use DomainKeys signing

No, do not use DomainKeys signing

Yes, use DomainKeys signing

Mail Job Specific Settings

Use inherited value: The user may change the setting supplied above

The user must use the setting supplied above without changes for each mail job

The user may change the setting supplied above on a per-job basis

Settings for Hosted LISTSERV Lists

Use inherited value: Do not enforce DomainKeys-signing

Enforce DomainKeys-signing for all LISTSERV Maestro standard postings to Hosted LISTSERV Lists.

Do not enforce this and allow the user to disable or enable DomainKeys-signing on a per-list basis.

To define the DomainKeys signature settings for a group, go to the Accounts and Identities screen, and then click on the name of the group. The Group Overview screen opens. Click on the **Group** menu, and then select **DomainKeys Settings**.

If your company or organization has decided that all messages that are sent using LISTSERV Maestro are to be signed with a DomainKeys signature, then choose **Yes, use DomainKeys signing** and **The user must use the setting supplied above without changes for each mail job**.

If your organization has decided that all messages sent from certain groups using LISTSERV Maestro are to be signed with a DomainKeys signature, then you can create a less strict policy by defining group-level deviations from the application-wide default settings. For example, one group may require the use of DomainKeys signatures while another group may not.

If it is not possible to agree on a strict policy for these settings even on the lowest group and/or single user level, then you should first choose a suitable default for enabling or disabling DomainKeys signatures, and then select the **The user may change the setting supplied above on a per-job basis** option.



Important: Before actually using such a setup, make sure to educate your users about the pros and cons of DomainKeys signatures. In a high volume environment, one reason to opt against DomainKeys signatures is that mail job delivery performance is impacted. LISTSERV uses highly optimized algorithms to perform

the signing and the throughput benefits from modern CPU extensions such as SSE2, so one option to use DomainKeys even in high volume environments is to acquire better hardware to run LISTSERV. On the other hand, not using DomainKeys may cause deliverability problems in the future if many of your subscribers have accounts with ISPs that enforce DomainKeys signatures on incoming mails or mark mails that lack such a signature as “coming from unsure origin” to warn the user about possible forgery or a phishing attempt.

The LISTSERV Maestro User Interface interacts with LISTSERV to determine if the supplied sender address is supported by one of the DomainKeys that were deployed to the LISTSERV host when DomainKeys was configured. This check is performed at several stages during the life cycle of an email job. The sender definition settings of an email job are only accepted as valid if either DomainKeys signing is switched off or if the check succeeds at the LISTSERV host that is configured for the account. After this, an additional check is performed when the email job is authorized for delivery. If the email job is configured for future delivery, then there is a considerable time window during which the administrator may opt to change the DomainKeys settings at the LISTSERV host. Therefore, if DomainKeys has been disabled during this time window, then the email job delivery will fail with an appropriate error message.

In addition, the system also performs consistency checks between the settings for the current email job and the settings that are defined for the current account in the Administration Hub. These checks ensure that the settings of the email job are the same as the default settings if the administrator has defined that users may not change the DomainKeys settings on the job level. Once the email job has been authorized for delivery, these additional checks are not repeated. This means that if the settings in the Administration Hub are changed, then this change will only affect email jobs that have not yet been authorized for delivery. Jobs that are authorized will not be affected from subsequent changes in the Administration Hub.

Section 22 LISTSERV & LISTSERV Maestro Integration

The goal of the LISTSERV and LISTSERV Maestro integration is to make the user experience a seamless experience when working with the LISTSERV Web Interface and with LISTSERV Maestro. This seamless integration will give the user's the perception that these two separate applications are actually working as one. The integration of LISTSERV and LISTSERV Maestro includes the following aspects:

- **LISTSERV and LISTSERV Maestro Interface Link** – This aspect deals with actually linking the LISTSERV Web Interface (WA) and the LISTSERV Maestro (LUI, HUB, and the membership areas) so that a menu appears within each interface, allowing users to switch between the two applications. For more information, see Section 22.1 [Defining the LISTSERV and LISTSERV Maestro Interface Links](#).
- **LISTSERV Web Interface (WA) and LISTSERV Maestro Single Sign-On** – This depends on the first aspect; the link between the WA and LISTSERV Maestro must be created for this aspect to work. If a link does exist, then this aspect deals with enabling the single sign-on feature, which allows users to switch between the two applications without having to log out and log back in. For more information, see Section 22.2 [Enabling Single Sign-On](#).
- **Membership Area as WA Subscriber's Corner** – This depends on the first aspect; the link between the WA and LISTSERV Maestro must be created for this aspect to work. If a link does exist, then this aspect deals with replacing the Subscriber's Corner with one or more Membership Areas. This would give the user the ability to switch from the Membership Area to the WA Archive Pages, or vice versa, with a single sign-on. For more information, see Section 22.3 [Linking the Membership Area and the Subscriber's Corner](#).

22.1 Defining the LISTSERV and LISTSERV Maestro Interface Links

For users who regularly work with both the LISTSERV Web Interface (WA) and LISTSERV Maestro (LUI and/or HUB), it is desirable to be able to switch from one to the other via direct menu items in each of the interfaces, so that you do not have to remember the URLs of each interface (or maintain your own bookmarks).

This requirement is fulfilled by defining a link between the user interfaces of LISTSERV (WA) and LISTSERV Maestro (LUI/HUB).

Once a link is defined, menu items will appear both in the WA, LUI, and the HUB, allowing the user of each interface to switch to certain pages in the other interface (however, without single sign-on yet, which requires additional configuration and is dealt with in Section 22.2 [Enabling Single Sign-On](#)).

To establish a link to the LISTSERV host of the LISTSERV Web Interface, click on the **Global Settings** menu, select **Maestro User Interface**, select **LISTSERV Web Interface Access**, and then finally select **LISTSERV Web Interface Links**. The LISTSERV Web Interface Links screen opens.

Figure 22-1 LISTSERV Web Interface Links

LISTSERV Web Interface Links

Linking to the Web Interface of a LISTSERV host enables the following features:

- In the linked LISTSERV Web Interface (WA): A "LISTSERV Maestro" menu is displayed for all administrators, with direct access links to this LISTSERV Maestro instance.
- In this LISTSERV Maestro instance: A "LISTSERV" menu with direct access links to the linked LISTSERV Web Interface is displayed for the "admin" account and for all non-admin accounts which use the linked LISTSERV host in their LISTSERV connection settings.
- The Data Administrators of this LISTSERV Maestro instance will have the option to also link Maestro datasets to the LISTSERV Web Interface (if the necessary right is granted). This in turn will have the effect that the datasets' member areas will contain direct access links to the LISTSERV archive pages and vice versa.

Click on the "edit" link to edit an existing link.
Use the "Create New Link" button to create a new link.

LISTSERV Host	LISTSERV Web Interface Access URL
	Edit

This screen displays a list of all interface links that are currently defined for this LISTSERV Maestro instance. Each existing interface link is displayed with the following values:

- **LISTSERV Host:** The host at which the linked LISTSERV instance is running.
- **LISTSERV Web Interface Access URL:** The URL under which the web interface of this LISTSERV instance is accessible (see below for details).



Tip: It is possible to link one and the same LISTSERV Maestro instance to the interfaces of several LISTSERV instances. This may seem strange at first glance but is logical because you can have different accounts and/or groups in LISTSERV Maestro that all may be connected to different LISTSERV instances, and of course you want to be able to create interface links to all these LISTSERV instances. In other words, you can create an interface link to any LISTSERV instance that is used by at least one of the accounts or groups in the Administration Hub (in the LISTSERV connection settings of the account or group), but each such LISTSERV instance can be linked only once.

Once an interface link to a certain LISTSERV instance is defined, the following will happen:

- **For a normal account, in LUI:** All accounts that use this LISTSERV instance will have a special **LISTSERV** menu in the main menu that allows them to directly access certain pages in the LISTSERV web interface of the linked LISTSERV instance (if the accessed LISTSERV web interface page is login protected, then the LISTSERV login information will have to be supplied after the menu item is clicked, unless single sign-on is also defined, see Section 22.2 [Enabling Single Sign-On](#)). Accounts that use a different LISTSERV instance will not see this menu, unless a link to their LISTSERV instance has also been defined.

- **For the admin account, in LUI and HUB:** The special **LISTSERV** menu in the main menu will include a sub-menu for this linked LISTSERV instance that allows the administrator to directly access certain pages in this LISTSERV's web interface (if the accessed LISTSERV web interface page is login protected, then the LISTSERV login information will have to be supplied after the menu item is clicked, unless single sign-on is also defined, see Section 22.2 [Enabling Single Sign-On](#)). Therefore, if there are several linked LISTSERV instances, this **LISTSERV** menu will have one entry for each of these linked instances. If, however, there are no linked LISTSERV instances at all, then the **LISTSERV** menu will be hidden.
- **In WA:** All accounts will have a special **Maestro** menu in the main menu that allows them to directly access certain pages of the LISTSERV Maestro instance from where the interface link was defined (the LISTSERV Maestro login information will have to be supplied after the menu item is clicked before the associated LISTSERV Maestro page is displayed, unless single sign-on is also defined, see Section 22.2 [Enabling Single Sign-On](#)).

To create a new link, click the **[Create New Link]** button.

The table lists all interface links that are currently defined. To edit an existing link, click on the **Edit** link associated with that link.

If you are editing a link, then the Edit LISTSERV Web Interface Link screen opens. If you are creating a new link, then the Define New LISTSERV Web Interface Link screen opens.

Figure 22-2 Editing an Existing Link

When defining a new interface link, select the LISTSERV host that you want to define a link for from the drop-down menu. This drop-down menu contains all of the LISTSERV hosts that are in use by any of the current user accounts or groups. Each host appears only once. Also, each host can only be linked once; therefore, the list contains only those hosts that have no link defined as of yet.

When editing an existing interface link, the LISTSERV host that this link points to is already defined and can no longer be changed (if the host is incorrect, simply delete this link and create a new one with the correct host).

For the selected host, you also need to provide the **Postmaster Address** and **Password** and the **TCPGUI-port** on which the host can be reached. When you select a host from

the drop-down menu, these values will be filled out automatically, taken from the LISTSERV connection settings of the first account or group that is found using this host. If necessary, you can change these values.

Finally, fill out the **Access URL** to the LISTSERV web interface of the LISTSERV instance that you are linking to. This has to be the same URL as you would enter into the address field of a browser if you wanted to access this LISTSERV web interface (WA) manually. This access URL usually has the following form:

`http://HOST:PORT/scripts/wa.exe`

where you replace HOST with the corresponding host name and PORT with the HTTP-port used on that server (if PORT is the default HTTP-port "80", then you can leave out the :PORT part so that your URL looks like this: `http://HOST/scripts/wa.exe`).



Note: Usually, the value for HOST is the same host name as the LISTSERV host defined at the top of the screen, but this is not necessarily true. For example, if the server has several host names or if the HTTP access is routed via a proxy, then the host name at the top of the screen must be the name by which the server can be reached on the TCPGUI-port, while the host name for the access URL must be the name by which the server can be reached via HTTP. In addition, sometimes the LISTSERV Web Interface is installed to use a different URL than the one described above; in this case, provide this URL instead.



Optionally, you can also enter the **LUI Access URL** that is used for accounts *without* an account mapping, when switching from WA to LUI. When a user switches over from WA to LUI, then the access URL for LUI is determined in the following way:

- If there is an account mapping for the user account, then the system will use the access URL of the LUI account that is mapped to the logged in WA account.
- If there is no account mapping for the user account, then the system will use the LUI access URL specified here on the LISTSERV Web Interface Link screen. If no URL is specified here, then the displayed default will be used.

To submit the settings, click the **[OK]** button; to exit without submitting, click the **[Cancel]** button.

To delete an existing interface link, click on the **Edit** link for that particular link, and then click the **[Delete Link]** button.

Some special considerations when working with several LISTSERV Maestro instances:

Each LISTSERV instance can only be linked to one single LISTSERV Maestro instance, i.e. if you happen to have several LISTSERV Maestro instances that all use the same LISTSERV instance, and you define an interface link to this LISTSERV in the Administration Hub of the first LISTSERV Maestro, and then you try to also define an interface link to the same LISTSERV in the Administration Hub of the second LISTSERV Maestro, then you will get an error message. This error message tells you that the given LISTSERV instance has already been linked by another LISTSERV Maestro instance, and includes an option for overriding this previous link with the new link. However, this override can cause some problems. If you should choose to override an existing interface link to a different LISTSERV Maestro instance, then this will have the following negative effect:

- In the second LISTSERV Maestro instance (the one for which you define the second interface link that now overrides the first link), you will get the expected **LISTSERV** access menu (for the affected accounts) that will also correctly send users to the web interface of the linked LISTSERV instance.
- In the LISTSERV Web Interface (WA) of the linked LISTSERV instance, you will similarly get the expected **Maestro** access menu that now will send all users to the second LISTSERV Maestro instance.
- The first LISTSERV Maestro instance, however, will be unaware of this change, i.e. in this instance, the interface link definition will still remain in place in the Administration Hub and LISTSERV Maestro will still show the **LISTSERV** access menus. Also, it will still allow users to switch from LISTSERV Maestro to the web interface of LISTSERV. However, if a user does this and switches from LISTSERV Maestro (of the first LISTSERV Maestro) to WA, and then the user tries to go back to LISTSERV Maestro with the menu provided in WA, then this menu will send the user to the second LISTSERV Maestro instance, not the first instance that the user came from.

Therefore, if you should choose to override an existing interface link (from a different LISTSERV Maestro instance), then you should not forget to also log in into the Administration Hub of this other LISTSERV Maestro instance and delete the interface link to the same LISTSERV instance.

22.2 Enabling Single Sign-On

The previous section describes how you can create a link between a LISTSERV Web Interface and the LISTSERV Maestro so that both contain menus that allow users to switch between the two interfaces. However, when using these menus, users will still be required to log in at the other interface manually, which can be quite cumbersome.

To avoid this, the single sign-on feature can be configured. This feature allows you to define that, if a user logs in to LISTSERV Maestro with a certain LISTSERV Maestro account, and then this user switches over to the WA, then the user will automatically be logged in at the WA with a certain LISTSERV account (and vice versa). For this, the following preconditions must be fulfilled:

- An interface link between LISTSERV and LISTSERV Maestro must have been defined, as described above.
- For a normal user, there must exist a LISTSERV Maestro account at the linked LISTSERV Maestro instance. The account must be configured to use the linked LISTSERV instance (so that the **LISTSERV** menu appears when the user logs in with this account). We call this the LMA-account below. (For the administrator, no additional preconditions must be fulfilled)
- For the user (either normal or admin), there must exist a LISTSERV account at the linked LISTSERV instance. This account takes the form of an email address for which a password must have been registered at the linked LISTSERV instance. We call this the WA-account below.

With these conditions fulfilled, you can now define the single sign-on feature for these two accounts, with the following effects:

- If the user logs in to LISTSERV Maestro with the LMA-account and then switches over to WA (with the **LISTSERV** menu), then the user will be automatically logged in

at the WA with the WA-account. (Although, the user may have to re-login at a later time once the automatically created login-ticket expires. This can be avoided by allowing the WA to store the login information in a cookie so that this re-login may happen automatically.)

- If the user logs in at the WA with the WA-account and then switches over to LISTSERV Maestro (with the **LISTSERV Maestro** menu), then the user will automatically be logged in at LISTSERV Maestro with the LMA-account.

Such a single sign-on for a pair of LMA- and WA-accounts is defined by creating a mapping between these two accounts. To establish a mapping between LISTSERV Maestro and the LISTSERV Web Interface, click on the **Global Settings** menu, select **Maestro User Interface**, select **LISTSERV Web Interface Access**, and then finally select **LISTSERV Web Interface Account Mappings**. The LISTSERV Web Interface Mappings screen opens.

Figure 22-3 LISTSERV Web Interface Mappings

LISTSERV Web Interface Mappings

Mapping a LISTSERV Maestro account to an account at LISTSERV will allow the account holder to switch from LISTSERV Maestro to the LISTSERV Web Interface and back without having to login again with each switch. If several LISTSERV Maestro accounts are combined in an identity, then instead this identity can be mapped to an account at LISTSERV, which will allow the account holders for each LISTSERV Maestro account in the identity to switch to the corresponding account at the LISTSERV that is defined for their account.

Note: This account mapping only works if also the account's LISTSERV host is "linked", i.e. if a link to the account's LISTSERV host has been defined on the [LISTSERV Web Interface Links](#) page. (The account's LISTSERV host is the host which is configured in the LISTSERV connection settings of the account.)

Click on the "edit" link to edit an existing mapping.
Use the "Create New Mapping" button to create a new mapping.

LISTSERV Maestro Account	Mapped LISTSERV Account	At LISTSERV Host	Host Is Linked	
jht - test1	jhubert@lsoft.com	guest.dc.lsoft.com	Yes	Edit
jht - test2	johannes.hubert@lsoft.com	guest.dc.lsoft.com	Yes	Edit
sample - francoise	francoise@jht.de	offler.lsoft-germany.de	Yes	Edit
sample - holly	holly@jht.de	offler.lsoft-germany.de	Yes	Edit
sample - jht	jht@jht.de	offler.lsoft-germany.de	Yes	Edit
Identity	Mapped LISTSERV Account	At LISTSERV Hosts	Host Is Linked	
FMB	francoise@lsoft.com	guest.dc.lsoft.com offler.lsoft-germany.de	Yes Yes	Edit

This screen displays a list of all account mappings that are currently defined. Each existing account mapping is displayed with the following values:

- **LISTSERV Maestro Account:** The LISTSERV Maestro account (the "LMA-account") that is part of the mapping (see above).
- **Mapped LISTSERV Account:** The LISTSERV account (the "WA-account") that is part of the mapping (see above), in form of an e-mail address.
- **At LISTSERV Host:** The LISTSERV host that is being used by the given LMA-account. For the single sign-on to work, the given WA-account must be a valid

account at this LISTSERV instance, i.e. there must be a password registered for the email address that constitutes this WA-account.

- **Host Is Linked:** Contains either the value “Yes” or “No”, depending on if an interface link has been defined for the given LISTSERV host or not. For the single sign-on to work, this value must appear as “Yes” (because otherwise there will not even be any access menu items either in LUI or the WA, which would leave the account mapping a moot point).

To create a new mapping, click the **[Create New Mapping]** button. The Define New LISTSERV Web Interface Account Mapping screen opens.

Figure 22-4 Creating a New Account Mapping

On this screen, select the LISTSERV Maestro user account (the “LMA-account”) from the drop-down list at the top. The list contains all unmapped standard accounts, as well as entries for all unmapped admin accounts (although there is only a single admin account, it can be mapped individually for each available LISTSERV instance so the list contains one admin account entry for each LISTSERV instance).

Next, you need to enter the email address that constitutes the LISTSERV account (the “WA-account”) that you want to map to.



Note: Each LMA-account can only be mapped once. In addition, you can only create one mapping for each WA-account (email address) on one and the same LISTSERV instance. In other words: If you have two LMA-accounts which use the same LISTSERV instance, then those two LMA-accounts must be mapped to two different WA-accounts (email addresses), since both WA-accounts are then interpreted as belonging to the same LISTSERV instance. If however you have two LMA-accounts which use different LISTSERV instances, then you can map those two LMA-accounts to one and the same WA-account (email address), since this one address is actually interpreted as two different accounts at two different LISTSERV instances, which only happen to use the same email address.

To edit or delete an existing account mapping, click on the **Edit** link in the matching table row. The editing of an existing account mapping is similar to what you do when creating a new account mapping (see above), except for that the selected LUI-account is now fixed.

Also when you edit an account mapping there is a **[Delete]** button that you can use to delete this account mapping.

To submit the settings, click the **[OK]** button; to exit the screen without submitting, click the **[Cancel]** button.

Some special considerations when working with “identities”:

The above describes the account mapping for normal LISTSERV Maestro accounts only, and how this enables single sign-on for accounts which are mapped.

There is, however, another topic here, in case you are using the Identity feature of LISTSERV Maestro. With identities, the following additional considerations apply:

- If a LMA-account is part of an identity, then you can no longer define a mapping for this account individually. Therefore, on the Define New Mapping screen, the drop-down menu that you can select the account to map will not contain this account.
- However, you can create a mapping for a whole identity. Therefore, if you have any identities defined, then, on the Define New Mapping screen, the drop-down menu that you can select the account to map does not only contain the available user accounts, but it also contains the available identities (those identities that are not already mapped).
- If a mapping is defined with an identity instead of an account, then the single sign-on works as follows instead:
 - If the user logs in to the LISTSERV Maestro with any of the LMA-accounts from the identity and then switches over to LISTSERV Web Interface (using the **LISTSERV** menu), then the user will be automatically logged in at the LISTSERV Web Interface with the LISTSERV Web Interface account mapped to the identity. (Although, the user may have to re-login at a later time if the automatically created login-ticket expires. This can be avoided by allowing the interface to store the login information in a cookie so that this re-login may happen automatically.)
 - If the user logs in at the LISTSERV Web Interface (WA) with the WA-account and then switches over to LISTSERV Maestro (using the **Maestro** menu), then the user will be presented with a selection page that shows all LMA-accounts in the mapped identity. Once one of the LMA-accounts has been selected from this list, then the user will automatically be logged in to LISTSERV Maestro with this LMA-account.
- It is allowed to combine LMA-accounts into an identity that does not use the same LISTSERV instance. Combined with the fact that an identity can only be mapped to only a single LISTSERV Web Interface account (email address), then the following situations may arise:
 - If a user logs in to LISTSERV Maestro with a LMA-account from the identity that uses a LISTSERV instance for which no interface link has been defined, then this user will not see the special **LISTSERV** access menu at all.
 - If a user logs in to LISTSERV Maestro with a LMA-account from the identity that uses a LISTSERV instance for which an interface link has actually been defined, but at this LISTSERV there exists no account that matches the mapped LISTSERV Maestro Web Interface account (i.e. there is no password

registered for this email address), then this user will see the special **LISTSERV** access menu. But, if the user clicks on any of its options, then they'll have to provide the login information at the LISTSERV Web Interface manually (if the user tries to access a protected page).

- If a user logs in to LISTSERV Maestro with a LMA-account from the identity that uses a LISTSERV instance for which an interface link as actually been defined, and at this LISTSERV there actually exists an account that matches the mapped LISTSERV Web Interface account (i.e. there is a password registered for this email address), then this user will see the special **LISTSERV** access menu. And, if the user clicks on any of its options, then they'll automatically be logged in at the LISTSERV Web Interface with the mapped LISTSERV Web Interface account. (Although, the user may have to re-login at a later time if the automatically created login-ticket expires. This can be avoided by allowing the interface to store the login information in a cookie so that this re-login may happen automatically.)

22.3 Linking the Membership Area and the Subscriber's Corner

In addition to the link between the LISTSERV and LISTSERV Maestro interfaces (as described above), it is also possible to link a dataset (or several datasets) with the LISTSERV Web Interface (WA) so that the membership areas of the linked datasets act as a replacement for the WA's normal subscriber's corner and that subscribers who login to a member area can access the archive pages of WA.

Such a link between a dataset and the WA is defined on dataset level, i.e. by the data administrator who has administrative access to the dataset in question. However, before the data administrator can define such a link for a given dataset, the following preconditions must be met:

- There must exist a normal interface link (as described in Section 22.1 [Defining the LISTSERV and LISTSERV Maestro Interface Links](#)) between the LISTSERV Maestro that contains the dataset and the LISTSERV instance that is configured for the data administrator who administrates the dataset (i.e. the LISTSERV instance configured in the LISTSERV connection settings of the data administrator's account or group).
- The data administrator must have been granted the additional user right to create links between datasets and the WA. To do this, go to the Accounts and Identities screen and click on the data administrator's account name. The User Account Overview screen opens. Click on the **User Account** menu, and then select **User Right Settings**. On the User Right Settings screen, check **The user may link Recipient Datasets to the LISTSERV Web Interface**.

Once these preconditions are fulfilled, the data administrator can define a link between a given dataset and the WA at any time using the dataset's definition wizard (see the [LISTSERV Maestro Data Administrator's Manual](#) for these instructions).



Note: Linking a dataset in this fashion has the additional effect that the **Member Password** option of the dataset is automatically set to **The member will get a**

system defined password, and this password can not be changed until the link to the WA is removed.

The link between a given dataset and the WA can be defined for one or more datasets and has the following effects:

- In the Membership Areas of the linked datasets there will appear two additional menu options that allow the subscribers to access the LISTSERV Archives and the Archive Search pages in the WA.
- In the WA, the menu options that point to the normal Subscriber's Corner will be hidden, and, in their place, there will be a menu that contains options to all Membership Areas that are linked to this WA.

So essentially, the linked member areas "take over" the subscriber's corner functionality of WA. The normal subscriber's corner will no longer be accessible, instead there will be links/menu-entries to the linked membership areas, which offer the subscriber the functionality that is normally offered by the subscriber's corner.

- These links and menu-entries between the membership areas and the WA archive-pages can be used by subscribers with "single sign-on" automatically enabled, i.e. a subscriber who logs into the membership area can switch to the WA archive-pages without having to login again. (Although, the subscriber may have to re-login at a later time, once the automatically created login-ticket expires. This can be avoided by allowing the WA to store the login information in a cookie so that this re-login may happen automatically.)

Similarly, if the subscriber first logs into the WA, then it is possible to switch to any of the linked membership areas without having to login again.

Additional Considerations:

As described above, it is possible to create a link between a dataset and the WA for several datasets at once. In this case, the WA will contain a menu that lists the Membership Areas of all linked datasets (by name), and the user can select the Membership Area to be directed to by selecting it from the menu. This means that, in the WA, the user will see the names of all linked datasets and will be able to switch to all of them, as long as the user is actually a member of the selected dataset.

As a result, when linking several datasets to the same WA, you need to carefully consider the datasets that you want to actually link, in order to avoid causing problems.

Consider the problems in the following situations:

- A football organization that uses LISTSERV Maestro to offer mailing lists for the fan clubs of various rivaling football teams. All datasets are administrated by one data administrator who is a member of the actual organization (not of one of the clubs).

Being sensible about the rivalries between the clubs, the data admin has of course created separate datasets for each of the clubs so that the fans of one team do not see the mailing lists dedicated to the other teams.

If the data administrator links several (or all) of these datasets to WA, then this separation would be broken on WA's side because, in the menu for the various

linked Membership Areas, there would appear the Membership Areas of all clubs, possibly offending some of the fans.

A better solution would be to create different LISTSERV Maestro groups and have each of the fan club datasets in a separate group, with a separate LISTSERV instance for each group. That way, the data administrator could link all fan club datasets to WA, as they would not be using the same WA.

- A similar situation, where however the data admin is not a member of the organization, and each fan club administrators its own datasets and mailing lists.

Because of this, the LISTSERV Maestro administrator has created separate LISTSERV Maestro groups, one for each fan club. In these groups, the administrator has created various accounts, one of which has the data administrator rights for that group, so that a member of each fan club can administrator the datasets and lists of that club. Therefore, there is a data administrator in each group, one for each fan club.

If all groups are connected to the same LISTSERV instance (via their LISTSERV connection settings in the Administration Hub), and the LISTSERV Maestro administrator grants the **The user may link Recipient Datasets to the LISTSERV Web Interface** user right to the data administrators of all groups, then it could happen that each data administrator decides independently to link his dataset to the WA. This would again have the effect that the links to various membership areas (of the various fan clubs) all appear in the same menu in the WA, which is definitely not a good idea.

In addition, since the various data administrators would not even know that the data administrators of other fan clubs have also connected their dataset (until they have a look at the WA menu), the data administrators would not even be aware of this (and even if they were aware of this, each of them would probably demand that the other data administrators remove their links).

Therefore, the LISTSERV Maestro administrator must take care to not simply grant the **The user may link Recipient Datasets to the LISTSERV Web Interface** user right to just any data administrator in order to avoid such conflicting situations.

A better solution would be to have separate LISTSERV instances for each of the groups (and fan clubs); in which case, it would then be no problem if all data administrators have this right, since they would all only affect their own WA.

Section 23 Moving a LISTSERV Maestro Installation

New

Sometimes it is necessary to replace the hardware on which a LISTSERV Maestro instance is installed, or to move a LISTSERV Maestro installation from one server to another. This section describes how to do this for the following LISTSERV Maestro installation types:

- All LISTSERV Maestro components are installed on the same server and the internal system database is used.
- All LISTSERV Maestro components are installed on the same server and an external system database is used, which is either also running on the same server, or is running on a different server.

Therefore, as long as LISTSERV Maestro is installed with all components on the same server (not counting the system database, which may also be on a different server), then this section describes how to replace or move this LISTSERV Maestro instance.

If however the LISTSERV Maestro components are distributed over several servers, and the components are to be moved separately, then please see [Section 24.2 Moving Components to another Server](#).

The exact procedure that you need to follow to replace or move a server depends on what exactly you want to accomplish and if the internal or an external system database is being used. The following subsections describe various possible scenarios. Find the scenario that matches your situation and follow the procedure described in that section.



Notes: The procedures only describe how to either replace the server hardware of an existing LISTSERV Maestro installation, or how to move an existing LISTSERV Maestro installation from one server to another.

It is **not** possible to **clone** an existing LISTSERV Maestro installation, i.e. to create a copy of an installation that contains the same data as the original installation, while having both installations running in parallel.

Therefore, the following scenario, for example, is **not possible**: Server A is running a LISTSERV Maestro installation. This installation is cloned to server B, so that server B now runs a LISTSERV Maestro installation with the same data. Server A continues to run too, for the time being, while the new installation on server B is being tested, to see if everything is working smoothly. Once this has been validated, server A is shut down (and uninstalled) and server B takes over its operation. This scenario is not only **impossible** because of technical reasons, but also because it does not make sense from an operational point of view either: During the time when both servers are running in parallel, server A would still be the "main" server, and server B would only be a copy that is currently undergoing verification. However, any data changes that happen on server A after the copying (while it is still the "main" server) would never be seen by server B. This includes jobs that are being created, delivered, or otherwise changed during this time, plus tracking events that arrive at server A during this time (or later). When the testing of server B would finally be complete and B would be ready to "take over" from server A, the data on B would

already be outdated in comparison to A, with no way to consolidate the two again. Therefore, it does not make sense to keep A running while the installation on B is being verified. But if A is shut down anyway at the moment the installation is moved to B, it is no longer a clone/copy operation, but instead either a replacement or move operation, as described below.

23.1 Server Replacement

Scenario: Replacing the LISTSERV Maestro Server with a New Server

The old LISTSERV Maestro server will be replaced with a new server. After the procedure, the LISTSERV Maestro installation will run on the new server, which will use the same host name as was previously used by the old server (possibly even the same IP-address), and the old server is either decommissioned or is used for other purposes, with a different host name (and possibly different IP-address).

This scenario, for example, applies if for some reason the old server hardware is getting old and a possible breakdown will be avoided, or the old hardware can no longer cope with the increased demands, so a new server with better hardware will replace it.



Note: With this scenario, LISTSERV Maestro will **never** be running in parallel on the two servers.

To perform such a server replacement, do the following:

1. On the old server, perform a backup, as described in Section 23.3 [Preparation for Replacing/Moving: Saving a Backup](#).



Important: If the backup fails with errors, then do not proceed! Resolve the errors first, then trigger another backup. Only proceed with a successful backup!

2. Immediately after the backup has completed successfully, shut down LISTSERV Maestro on the old server.
3. **Do not start the old LISTSERV Maestro installation again while performing the rest of the procedure!** Remember that LISTSERV Maestro can also be restarted automatically as a service (Windows) or daemon process (Unix/Linux), if the server computer is restarted. This must also be avoided!
4. If you are planning to still use the old server after the procedure, for other purposes and with a new hostname (and possibly IP-address), then change the hostname/IP-address of the old server now.

If you are planning to decommission the old server and not use it anymore, then you still need to temporarily change its hostname (and possibly IP-address) at this point, to a temporary hostname/IP-address.



Note: Changing the hostname may also require a DNS-mapping entry for the new hostname.

5. The old hostname is now no longer used by the old server and is therefore "free" for the new server.

Configure the new server to use the hostname (and possibly IP-address) that previously was used by the old server. Also, check the DNS-mapping of that hostname to make sure that it points to the IP-address that is now used by the new server.

6. If LISTSERV Maestro is configured to use an external database, and this external database was also installed on the old server, then you need to install a new external database on the new server too.

This step is **not** necessary if LISTSERV Maestro uses the internal database or if it uses an external database on some altogether different server (not the old server). In these cases, skip to the next step.

But, if the external database was used **and** it was installed on the old server, then now you need to install the same database software also on the new server. In this database software, create a new system database (or database schema), with the same name, user account (including password and access rights), and other configurations as was previously used by the system database on the old server. If necessary, check the configuration of the system database on the old server. (See Section 10.2 [Preparing the System Database](#).)

7. Install LISTSERV Maestro on the new server. Install either the same or a newer version than was installed on the old server. (Do not install an older version/build!)

If possible, install into a folder with the same name and path as on the old server (on Windows also on the same drive). This is not strictly necessary, but makes the rest of the procedure easier.

8. If the old LISTSERV Maestro installation also included a LISTSERV installation on the old server, then you need to include LISTSERV in the installation on the new server too. In this case, you also need to transfer all the LISTSERV data from the old server to the new server (see the LISTSERV manuals for how to do this).
9. If you want to retain the log file history (recommended) then copy the log files from the old server to the new server. Log files are stored in three different folders. Copy them from each folder on the old server to the corresponding folder on the new server:

```
[maestro_install_folder]/hub/logs
```

```
[maestro_install_folder]/lui/logs
```

```
[maestro_install_folder]/trk/logs
```

10. If you have added any additional content to the webapps folder of the installation on the old server, as described in Section 26 [Adding Content to the Tomcat Server](#), then copy this content to the webapps folder of the installation on the new server.

Similarly, if you have any other custom content or custom folders in the LISTSERV Maestro installation folder on the old server, then copy it to suitable locations on the new server.

11. Locate the backup that you saved in step 1 (on the old server) and restore it into the fresh LISTSERV Maestro installation on the new server. For details, See Section 11.7 [Restoring a Backup](#).

Restoring the backup also includes the step of starting LISTSERV Maestro on the new server.

12. If the new LISTSERV Maestro instance on the new server was installed into a different folder than on the old server (or on Windows onto a different drive), then login to

the Administration Hub and check all settings that may contain folder paths, and if necessary change them to the correct path/drive-letter for the new installation location (changes do not require a restart):

- **Backup Folders:** Under **Global Settings > Administration Hub/Maestro User Interface/Maestro Tracker > General Administration** (check for each of the three components!).
 - **Content Upload Restrictions:** Under **Global Settings > Maestro User Interface > Default Content Restrictions**. Also, check for any overrides on account or group level under **User Account/Group > Content Restrictions** (for all non-group accounts and all groups).
 - **Recipient Upload Restrictions:** Under **Global Settings > Maestro User Interface > Default Recipients Restrictions**. Also check for any overrides on account or group level under **User Account/Group > Recipients Restrictions** (for all non-group accounts and all groups).
 - **Hosted Data Import Restrictions:** Under **Global Settings > Maestro User Interface > Default Hosted Data Settings**. Also check for any overrides on account or group level under **User Account/Group > Hosted Data Settings** (for all non-group accounts and all groups).
13. Once LISTSERV Maestro is running properly on the new server, with the backup restored, login to the Administration Hub on the new server and go to **Global Settings > Maestro User Interface > General Administration** and unselect the **Outbox: Sending is disabled** and **Login Access: LISTSERV Maestro User Interface is locked** checkboxes. Click **[OK]** to submit your changes.



Important: Once you have verified that LISTSERV Maestro is running properly on the new server, make sure to uninstall the old LISTSERV Maestro installation on the old server.

If desired, you can first make a backup of the LISTSERV Maestro installation folder on the old server before you uninstall it (by doing a simple file system copy of the whole installation folder), but you must uninstall it in any case.

14. You can now either decommission the old server or use it for other purposes (with its new hostname).

23.2 Moving a LISTSERV Maestro Installation

Scenario: Moving LISTSERV Maestro Installation From Old Server to New Server

The LISTSERV Maestro installation will be **moved** from the old server to a new server. The new server will use a different host name (and IP-address) than the old server and the old server will continue to use its original hostname (and IP-address).

This scenario, for example, applies if for some reason the installation must be moved to a new server, but for other purposes, the old server must remain operational with the same host name as before, only now without a LISTSERV Maestro installation on it.



Note: With this scenario, LISTSERV Maestro will **never** be running in parallel on the two servers.



Important: Moving a LISTSERV Maestro installation to a server with a **different** hostname is **not a good idea** and should be avoided, unless absolutely necessary.

The reason for this is the following: When LISTSERV Maestro sends out e-mails with tracked links, then the URLs of these tracked links point back to the LISTSERV Maestro server, using the server's hostname. But, if LISTSERV Maestro's hostname is changed (which would be the case if the LISTSERV Maestro installation is moved to a server with a different name), then all the already delivered e-mails still contain tracking links that point to the old hostname. Only now there is no LISTSERV Maestro installation with that hostname anymore. This has an effect both on the recipients of the already delivered messages and on the LISTSERV Maestro user who owns the mail jobs for these messages:

- **Negative effect for the recipients** – All tracked links in all e-mails that were delivered before the server was moved cease to function. If a recipient clicks one of these links, then instead of seeing the intended target page, he will see a "server not found" message (or similar).
- **Negative effect for the LISTSERV Maestro user** – Since the tracked links cease to function, no tracking events for these links will reach LISTSERV Maestro. Therefore, all mail jobs that were delivered before the server was moved will no longer receive any new tracking events. Their tracking data will appear "frozen" in the state that it had at the moment the server was moved.

Mail jobs that are delivered after the server was moved are not affected by this.

Therefore, unless the negative effects described above are acceptable in your situation, you should never move a LISTSERV Maestro installation to a server with a new name. If you have to move to a new server, then consider if it is not possible to also "move" the hostname from the old server to the new server. I.e. assign a new hostname to the old server and assign the old hostname to the new server (possibly as an alias, if the new server already had a different hostname before). If you can find a way to also "move" the hostname to the new server, then you are no longer dealing with "moving" a LISTSERV Maestro installation, as described in this section, but instead you are dealing with a "server replacement". In this case, please see Section 23.1 [Server Replacement](#) for details.

But if you have decided that you do indeed want to move a LISTSERV Maestro installation to a new server, with a new hostname, in spite of the negative effects that this has, then follow these steps:

1. On the old server, perform a backup, as described in Section 23.3 [Preparation for Replacing/Moving: Saving a Backup](#).



Important: If the backup fails with errors, then do not proceed! Resolve the errors first, then trigger another backup. Only proceed with a successful backup!

2. Immediately after the backup has completed successfully, shut down LISTSERV Maestro on the old server.
3. Do not start the old LISTSERV Maestro installation again while performing the rest of the procedure! Remember that LISTSERV Maestro can also be restarted auto-

matically as a service (Windows) or daemon process (Unix/Linux), if the server computer is restarted. This must also be avoided!

4. Install LISTSERV Maestro on the new server. Install either the same or a newer version than was installed on the old server. (Do not install an older version/build!)

If possible, install into a folder with the same name and path as on the old server (on Windows also on the same drive). This is not strictly necessary, but makes the rest of the procedure easier.

5. If you want to retain the log file history (recommended) then copy the log files from the old server to the new server. Log files are stored in three different folders. Copy them from each folder on the old server to the corresponding folder on the new server:

```
[maestro_install_folder]/hub/logs
```

```
[maestro_install_folder]/lui/logs
```

```
[maestro_install_folder]/trk/logs
```

6. If you have added any additional content to the webapps folder of the installation on the old server, as described in Section 26 [Adding Content to the Tomcat Server](#), then copy this content to the webapps folder of the installation on the new server.

Similarly, if you have any other custom content or custom folders in the LISTSERV Maestro installation folder on the old server, then copy it to suitable locations on the new server.

7. Locate the backup that you saved in step 1 (on the old server) and restore it into the fresh LISTSERV Maestro installation on the new server. For details, See Section 11.7 [Restoring a Backup](#).

Restoring the backup also includes the step of starting LISTSERV Maestro on the new server.

8. Since you have restored the backup on a server with a different hostname, make sure to go through all configuration settings and adjust them to use the new hostname.

- Check the INI-files of the three components (changes require a LISTSERV Maestro restart):

```
[maestro_install_folder]/lui/lui.ini
```

```
[maestro_install_folder]/hub/hub.ini
```

```
[maestro_install_folder]/trk/tracker.ini
```

- Login to the Administration Hub and check the URL settings (changes do not require a restart) under **Global Settings > Maestro User Interface > Default URL Settings**. Also check for any overrides on account or group level under **User Account/Group > URL Settings** (for all non-group accounts and all groups).

9. If the new LISTSERV Maestro instance on the new server was installed into a different folder than on the old server (or on Windows onto a different drive), then login to the Administration Hub and check all settings that may contain folder paths, and if

necessary, change them to the correct path/drive-letter for the new installation location (changes do not require a restart):

- **Backup Folders:** Under **Global Settings > Administration Hub/Maestro User Interface/Maestro Tracker > General Administration** (check for each of the three components!).
 - **Content Upload Restrictions:** Under **Global Settings > Maestro User Interface > Default Content Restrictions**. Also, check for any overrides on account or group level under **User Account/Group > Content Restrictions** (for all non-group accounts and all groups).
 - **Recipient Upload Restrictions:** Under **Global Settings > Maestro User Interface > Default Recipients Restrictions**. Also check for any overrides on account or group level under **User Account/Group > Recipients Restrictions** (for all non-group accounts and all groups).
 - **Hosted Data Import Restrictions:** Under **Global Settings > Maestro User Interface > Default Hosted Data Settings**. Also check for any overrides on account or group level under **User Account/Group > Hosted Data Settings** (for all non-group accounts and all groups).
10. Once LISTSERV Maestro is running properly on the new server, with the backup restored, login to the Administration Hub on the new server and go to **Global Settings > Maestro User Interface > General Administration** and unselect the **Outbox: Sending is disabled** and **Login Access: LISTSERV Maestro User Interface is locked** checkboxes. Click **[OK]** to submit your changes.



Important: Once you have verified that LISTSERV Maestro is running properly on the new server, make sure to uninstall the old LISTSERV Maestro installation on the old server.

If desired, you can first make a backup of the LISTSERV Maestro installation folder on the old server before you uninstall it (by doing a simple file system copy of the whole installation folder), but you must uninstall it in any case.

11. You can now continue to use the old server for other purposes (with its old host-name).

23.3 Preparation for Replacing/Moving: Saving a Backup

The replacement/moving procedures described in the previous sections require you to perform a safe backup as their first step. The procedure on how to do this safe backup is described here:

1. Login to the Administration Hub with the admin account.
2. Go to **Global Settings > Maestro User Interface > General Administration**.
3. Select the **Outbox: Sending is disabled** checkbox.
4. Select the **Login Access: LISTSERV Maestro User Interface is locked** checkbox.
5. Click **[OK]**.
6. Click on the **Switch to LUI** link at the top of the screen to go to User Interface.
7. Go to **Mail Job > Mail Jobs Overview**.

8. Select the **Jobs In System** tab.
9. Click on the **Display the jobs owned by** drop-down menu and select **<Any Owner>**.
10. Select the **Ongoing Jobs** sub-tab.
11. Verify that the job list on this sub-tab is either empty, or if there are any jobs remaining, that none of them are "**being processed**". The "**being processed**" state is shown with the **yellow-arrows icon** in the **State** column at the right. Any job that is still in the list must either be in the "**authorized**" state (**clock icon**) or "**failed**" state (**red-X icon**).

If there are still jobs in the "**being processed**" state, then wait for these jobs to be completed: If they are completed successfully, they will be removed from this list. If they fail, they will change into the "**failed**" state on this list.

While there are still jobs "**being processed**", refresh the list regularly until all jobs are either gone or only "**failed**" and "**authorized**" jobs remain.
12. Once there are no jobs "**being processed**" anymore, use the **Switch to HUB** link at the top of the screen to go back to the HUB.
13. Go to **Global Settings > Administration Hub > General Administration**.
14. Click the **[Execute Backup Now]** button.
15. Logout.
16. Now you need to wait for the backup to complete. To find out if the backup is complete, monitor the most recent log file of the Administration Hub. This file is located in the following folder:
`[maestro_install_folder]/hub/logs`

The name of the log file is similar to "HUB-YYYYMMDD.log" where YYYYMMDD is the date of the current day (in order "year month day").

Open this log file and check the most recent entries at the end of the log. You should see entries which say that a backup was started and that the backup is running.

Once the backup is complete, you will see a message like the following (with a matching date and time):

```
Final completion date: Fri Apr 16 00:01:01 CEST 2010  
The backup was completed successfully
```

You may have to reopen/refresh the log file several times until you see this message appear. Be patient - especially if you have a larger LISTSERV Maestro installation, the backup may take a while.
17. Once you see the above message you know that the backup has completed successfully. The backup will have been stored in the most recent backup folders. There is a backup folder for each of the components LUI, HUB, and TRK, so the full backup consists of the contents of all three backup folders of these components.
18. You now have performed a successful "safe backup" and can go back to the replacement/moving procedure that you were originally executing.

Section 24 Distributed Components

The three LISTSERV Maestro components, the LISTSERV server(s), the SMTP server(s), and the optional external database(s) may be installed on any combination of hosts, from one single host shared by all components to six or more dedicated hosts, one for each component. If different components are installed on separate servers, it is not necessary that all of the servers have the same operating system. It is possible to install the Maestro User Interface and Administration Hub components on a Windows server and at the same time the Maestro Tracker component on a Linux server and LISTSERV on Solaris (or other combinations). For more information on host restrictions, installing LISTSERV Maestro, and starting and stopping the LISTSERV Maestro service, see the LISTSERV Maestro Installation Manual.

Distributing components has several advantages:

- **Load Distribution** – Processor and disk load is shared between several servers, giving each component more “room” to operate.
- **SSL Security** – If SSL secure access (that is, HTTPS) is required for HUB and/or LUI, those must be on a separate server than TRK, which must not be configured to use SSL.
- **Separate Maintenance** – Not all components have to be shut down or re-started whenever a maintenance task on one of them requires it. The other components may continue running (although when a component that other components rely on is shut down, the others may have to wait for that component to come back up before they can finish performing any tasks).

Specifically, the Maestro Tracker component has very rigid uptime requirements. This component should constantly be running to be able to collect the tracking data from the messages that are sent. It can only do so while it is running and connected to the Internet. Therefore, it is not a good idea to shut down the server on which the Maestro Tracker component is running – this should only be done as a last resort. Other components do not have these strict uptime requirements. To minimize Tracker downtime in the event of maintenance on the tracker components, it is a good idea to have the Maestro Tracker component on a separate server.

For optimal performance for a high-volume installation with requirements for high availability and high performance, a component distribution on five or more servers might be necessary:

- **User-Interface and Hub Server** – Contains the LISTSERV Maestro components Maestro User Interface and Administration Hub.
- **Tracker Server** – Contains the Maestro Tracker component.
- **Database Server** – Contains the Maestro System Database component.
- **LISTSERV Server** – Contains the LISTSERV external component.
- **SMTP Server** – Contains the SMTP service.
- Depending on load and performance requirements, you may need additional LISTSERV servers to assign different LISTSERV instances to each account group

and/or for bounce processing and/or to serve as “distributed workers” for the primary LISTSERV server(s). To handle high volumes of deliveries quickly, you may also need additional SMTP servers.



Tip: L-Soft Consulting Services can assist you in finding the right configuration for your needs. Contact your L-Soft sales representative for more information.

24.1 Fresh Installation with Distributed Components

A fresh installation with distributed components is a straightforward operation. To install any of the three LISTSERV Maestro components, simply run the LISTSERV Maestro setup on the server where the component(s) will be installed and then select the required components from the list, while leaving all components to be installed on other servers unchecked. The other external components (LISTSERV, external database, and SMTP server) are installed separately. Simply execute each application setup on the respective server(s).

24.2 Moving Components to another Server

If you already have a running installation and plan to move one or several of its components to other servers, you can do so too (this also applies, in slightly different form, if you need to change the host name of the server where your components are installed).

The following sections describe for each component, what you need to do to move it to a different server.

24.2.1 Moving the Maestro User Interface Component to Another Server

Follow these steps to move the Maestro User Interface Component to another server:

1. As a preparation for moving the Maestro User Interface component, first think about which kind of system database you are using:
 - a. If you are currently using an **external system database**, then you need to make sure that the new server will also be able to contact this external database over the network. You must also remember to add the database driver file(s) for that external database to the fresh installation of your Maestro User Interface on the new server (see step 4).
 - b. If you are currently using the **internal system database**, then you must decide whether or not you want the new server to continue using the internal system database or if you want to switch to an external system database instead.

If you decide to stay with the internal system database, then simply continue with the procedure described here.

If you decide to switch to an external database, then, for the time being, abandon the procedure described here and *first* switch to the new database, as described in Section 10.1 [Configuring the External System Database](#). Then return to here, only that now you proceed as described in the bullet above, which deals with an external system database.

2. Only required if you are using the internal system database: Trigger a backup of LISTSERV Maestro. In the Administration Hub, click on the **Global Settings** menu, select **Administration Hub**, and then finally select **General Administration**. The General Component Settings for Administration Hub screen opens; click on the **[Execute Backup Now]** button. Wait for the backup to complete. Check the backup log in the “hub/logs” folder in the LISTSERV Maestro installation folder on the server where the Administration Hub is installed. A backup log with the current date and time must appear, and the log must state that the backup has completed successfully.
3. Shut down the existing LISTSERV Maestro installation on all servers where components of it are installed. If the internal database is in use and Hosted LISTSERV Lists are being used, then shut down all LISTSERV instances as well.
4. The Maestro User Interface must not already be installed on the new server. However, the Administration Hub or Maestro Tracker components may already be installed, in which case the Maestro User Interface is added to the existing installation:
 - a. If there are no LISTSERV Maestro components already installed on the new server, then simply execute a fresh installation as described in the installation manual for your operating system. During the installation, when queried for which components to install, select only the Maestro User Interface component.
 - b. If there are already other LISTSERV Maestro components installed on the new server, then start the installation package for your operating system in the same way as if doing a fresh installation. The installation package will then recognize the existing installation and will give you the option of adding new components to it. Select the Maestro User Interface component to be added and proceed with the installation.

In both cases, if you are planning to use the internal database on the new server, then remember to also select the MySQL component during installation. If you are not using the internal database, then do not select the MySQL component for installation on the new server.

After the installation, do *not* start LISTSERV Maestro!

5. Transfer the following files and folders (including all files and subfolders in them) from the previous server to the new server:

```
[maestro_install_folder]/lui/lui.ini  
[maestro_install_folder]/lui/luidata  
[maestro_install_folder]/lui/registry
```

To transfer these files and folders, first *delete* them on the new server, then replace them with copies of the corresponding files and folders from the previous server (do not just copy the folders from the previous server over the same folders on the new server, since this may result in an inconsistent mix of files from the two servers).

(Depending on your installation, these paths may be slightly different on one or both of your servers.)

6. Only required if you are using the internal system database: Copy the backup folder into which the Maestro User Interface (LUI) backup was written during step above. By default, the folder will have a path like:

```
[maestro_install_folder]/lui/backup
```

However, this name may differ if you have configured a different backup folder for the Maestro User Interface component in the Administration Hub.

Locate the backup folder and copy it to the new server (to a temporary folder outside of the LISTSERV Maestro installation structure).


Then, edit the following file:

```
[maestro_install_folder]/lui/lui.ini
```

Add an entry like the following:

```
RestoreBackup=PATH_TO_BACKUP_FOLDER
```

where you replace “PATH_TO_BACKUP_FOLDER” with the path of the backup folder which you have just copied to this new server (see above). You can either give an absolute path or a path relative to the location of the “lui.ini” file.

7. On the old server, you now need to remove the previous installation of the Maestro User Interface component:
 - a. On Windows: Use Windows’ **Add/Remove Programs** panel on the old server to start the maintenance setup of LISTSERV Maestro. In the setup, choose **Modify**, and then deselect the **Maestro User Interface** component so that it is uninstalled.
 - b. On Linux / Solaris: Start the installation package for your operating system in the same way as if doing a fresh installation (see installation manual). The installation package will then recognize the existing installation and will give you the option of removing components from it. Select the Maestro User Interface component to be removed and proceed with the uninstallation.
8. Only required if you are using the internal system database: If Hosted LISTSERV Lists are being used, then the database access configuration of all connected LISTSERV instances needs to also be changed so that it points to the internal database on the new server. See the installation manual for details.
9. Restart LISTSERV Maestro (on all servers), as well as any LISTSERV instances that have been stopped.
10.  After LISTSERV Maestro has been started again, check the lui.ini settings and the configuration settings in the HUB to see if any need to be adjusted to the new server (for example, if the new server has a different host name than the old server).

24.2.2 Moving the Administration Hub Component to Another Server

Follow these steps to move the Administration Hub to another server:

1. Shut down the existing LISTSERV Maestro installation on all servers where components of it are installed. If the internal database is in use and Hosted LISTSERV Lists are being used, then shut down all LISTSERV instances as well.
2. The Administration Hub must not already be installed on the new server. However, the Maestro User Interface or Maestro Tracker components may already be installed, in which case the Administration Hub is added to the existing installation:
 - a. If there are no LISTSERV Maestro components already installed on the new server, then simply execute a fresh installation as described in the installation manual for your operating system. During the installation, when queried for which components to install, select only the Administration Hub component.
 - b. If there already are other LISTSERV Maestro components installed on the new server, then start the installation package for your operating system in the same way as if doing a fresh installation. The installation package will then recognize the existing installation and will give you the option of adding new components to it. Select the Administration Hub component to be added and proceed with the installation.

After the installation, do *not* start LISTSERV Maestro!

3. Transfer the following files and folders (including all files and subfolders in them) from the previous server to the new server:

```
[maestro_install_folder]/hub/hub.ini
```

```
[maestro_install_folder]/hub/accountreg
```

```
[maestro_install_folder]/hub/hubreg
```

To transfer these files and folders, first delete them on the new server, then replace them with copies of the corresponding files and folders from the previous server (do not just copy the folders from the previous server over the same folders on the new server, since this may result in an inconsistent mix of files from the two servers).

(Depending on your installation, these paths may be slightly different on one or both of your servers.)

4. Edit the following file of the Maestro User Interface component, which may be installed on a different server:

```
[maestro_install_folder]/lui/lui.ini
```

In the file, you need to edit the “RegistryHubHost” entry so that it contains the host name of the new server where the Administration Hub will be running.

5. Edit the following file of the Maestro Tracker component, which may also be installed on a different server:

```
[maestro_install_folder]/trk/tracker.ini
```

In the file, you need to edit the “RegistryHubHost” entry so that it contains the host name of the new server where the Administration Hub will be running.

6. On the old server, you now need to remove the previous installation of the Administration Hub component:
 - a. On Windows: Use Windows' **Add/Remove Programs** panel on the old server to start the maintenance setup of LISTSERV Maestro. In the setup, choose **Modify**, and then deselect the **Administration Hub** component so that it is uninstalled.
 - b. On Linux / Solaris: Start the installation package for your operating system in the same way as if doing a fresh installation (see installation manual). The installation package will then recognize the existing installation and will give you the option of removing components from it. Select the Administration Hub component to be removed and proceed with the uninstallation.
7. Restart LISTSERV Maestro (on all servers), as well as any LISTSERV instances that have been stopped.
8. After LISTSERV Maestro has been started again, check the configuration settings in the HUB to see if any need to be adjusted to the new server (for example, if the new server has a different host name than the old server).



24.2.3 Moving the Maestro Tracker Component to Another Server

Moving the Maestro Tracker component to a different server must be thoroughly planned. The problem is that all mails that were sent while the Maestro Tracker component was still installed on the old server will include mail tracking code with the old server name. If you now shut down and uninstall the Maestro Tracker component on that old server, then all future tracking events from those mails will be lost.

Even worse, click-through tracking links will not work any longer at all. This means that if a recipient clicks on a click-through tracked link that is connected to the old Maestro Tracker component's host name, then the recipient will receive a "Host not found" or "Page not found" error (it will look like a broken link – instead of being routed to the actual link target).

Because of this, you should be very careful when moving the Maestro Tracker component to a different server. Under normal production conditions, this should never be done. Only if your last tracked mailing is already some time in the past and you don't care about tracking events that get lost and broken links (because since already some time has passed since the mailing, there are no tracking events being generated any longer anyway, or at least only very few).

Of course, you can also solve this problem by changing the DNS registration of the host name. If the host name for your Maestro Tracker component was previously DNS-registered to point to the IP-address of the old server, then you can change the registration and let it point to the address of the new server instead. From the outside, this will look like there was no change at all (keep in mind that the propagation of a DNS change always takes a few days so, in the interim period, the adverse effects of moving the component, as described above, may still happen).

However, if you really need to or want to move your Maestro Tracker component, do so here:

1. Shut down the existing LISTSERV Maestro installation on all servers where components of it are installed. If the internal database is in use and Hosted LISTSERV Lists are being used, then shut down all LISTSERV instances as well.

2. Maestro Tracker must not already be installed on the new server. However, the Maestro User Interface or Administration Hub components may already be installed, in which case Maestro Tracker is added to the existing installation:
 - a. If there are no LISTSERV Maestro components already installed on the new server, then simply execute a fresh installation as described in the installation manual for your operating system. During the installation, when queried for which components to install, select only the Maestro Tracker component.
 - b. If there are already other LISTSERV Maestro components installed on the new server, then start the installation package for your operating system in the same way as if doing a fresh installation. The installation package will then recognize the existing installation and will give you the option of adding new components to it. Select the Maestro Tracker component to be added and proceed with the installation. After the installation, do not start LISTSERV Maestro!

3. Transfer the following files and folders (including all files and subfolders in them) from the previous server to the new server:

```
[maestro_install_folder]/trk/tracker.ini  
[maestro_install_folder]/trk/data
```

To transfer these files and folders, first delete them on the new server, then replace them with copies of the corresponding files and folders from the previous server (do not just copy the folders from the previous server over the same folders on the new server, since this may result in an inconsistent mix of files from the two servers).

(Depending on your installation, these paths may be slightly different on one or both of your servers.)

4. Edit the following file of your Maestro User Interface component (which may be installed on an entirely different server):

```
[maestro_install_folder]/lui/lui.ini
```

In the file, you need to edit the “TrackerHost” entry so that it contains the host name of the new server where Maestro Tracker will be running (of course, if you also changed the DNS-registration of the old host name to now point to the new server, then you do not have to do this change, since the actual host name is not changed – it only points to a different server).

5. On the old server, you now need to remove the previous installation of the Maestro Tracker component:
 - a. On Windows: Use Windows’ **Add/Remove Programs** panel on the old server to start the maintenance setup of LISTSERV Maestro. In the setup, choose **Modify** and deselect the **Maestro Tracker** component, so that it is uninstalled.
 - b. On Linux / Solaris: Start the installation package for your operating system in the same way as if doing a fresh installation (see installation manual). The installation package will then recognize the existing installation and will give you the option of removing components from it. Select the Maestro Tracker component to be removed and proceed with the uninstallation.



- Restart LISTSERV Maestro (on all servers), as well as any LISTSERV instances that have been stopped.
- If the new server of the Maestro Tracker component is now known with a different name than the previous server was, you also need to adjust the **Tracking URL** setting in the Administration Hub. For this, log into the Administration Hub and set the correct tracking URL either on global default level, group level or user level, whichever is applicable.

24.2.4 Moving the Database External Component to Another Server

- Install the database software on the new server and start it.
- Follow the instructions in Section 10.1 [Configuring the External System Database](#); however, instead of moving from one type of database to a different one, you move between two database which are on different servers, but may actually be the same type of database (e.g. database vendor, version, etc.).
- If you have done any special configurations or optimizations to your original database (with configuration tools or by editing configuration or INI files), then remember to apply the same adjustments to the new database installation too.

24.3 Server Name Aliases and Proxies


With any given installation of LISTSERV Maestro, the components of LISTSERV Maestro are installed on one or more servers, where each server has its own host name. Components on separate servers use the other server or servers' name(s) to access the component(s) there. Similarly, the "outside" world (users and email messages that are being tracked) accesses the components with their server names as well.

In the simplest setup, each server hosting a LISTSERV Maestro component will have a DNS name that can be used both for the inter-component communication as well as for "outside" world access. In this case, setup is straightforward and no extra measures have to be taken.

However, there are configurations in which the host names of the LISTSERV Maestro component servers are names known only in the local network, with no DNS names assigned. Or the hosts are, for security reasons, not accessible directly from the outside. Instead, there is a proxy (or other kind of "forwarder") that sits between the local network and the outside world so that the outside only ever knows the host name (and IP address) of the proxy, but never the names and addresses of the servers behind it (which also may be addresses from a local range, like the 192.168.0.0 subnet).

The figure below shows such a setup, where only the proxy has a valid non-local IP address and a registered DNS name (or several names, see examples following the figure), while the LISTSERV Maestro servers have only local names and addresses.

Figure 24-1 Sample Proxy Setup



Example 1

Assume that the proxy has a single DNS name “maestro.sample.com”. It could be configured to:

- Forward access on
maestro.sample.com:9001
to local host LUI (192.168.1.1), port 80
- Forward access on
maestro.sample.com:9002
to local host HUB (192.168.1.2), port 8080
- Forward access on
maestro.sample.com:9003
to local host TRK (192.168.1.3), port 8888

This example shows how a single DNS name can be “split” to proxy for three different servers, by employing different ports (9001-9003), which are mapped to different hosts (LUI, HUB, TRK) and their corresponding ports (80, 8080, 8888). Users wanting to access the Maestro User Interface would have to use a URL similar to: `http://maestro.sample.com:9001/loi`.

Users accessing the Administration Hub would use: `http://maestro.sample.com:9002/hub`

The tracking URLs would contain the URL `http://maestro.sample.com:9003/trk`

Example 2

As a second example, assume that the proxy has three assigned DNS names `lui.sample.com`, `hub.sample.com` and `trk.sample.com`, which are used to decide which local host to access, so the proxy could be configured to do the following:

- Forward access on
`lui.sample.com:80`
to local host LUI (192.168.1.1), port 80
- Forward access on
`hub.sample.com:80`
to local host HUB (192.168.1.2), port 8080
- Forward access on
`trk.sample.com:80`
to local host TRK (192.168.1.3), port 8888

In this example the “splitting” is realized by using three different host names, all assigned to the same server, where access on the standard HTTP-port 80 is mapped to the different local hosts (LUI, HUB, TRK) and their corresponding ports (80, 8080, 8888) depending on the DNS name used to access the proxy. Users wanting to access the Maestro User Interface would have to use a URL like `http://lui.sample.com/lui`. Users accessing the Administration Hub would use `http://hub.sample.com/hub` and the tracking URLs would contain the URL `http://trk.sample.com/trk`.

This example demonstrates that the host names of the servers hosting the LISTSERV Maestro components may differ when viewed locally or from the “outside” world. Internally, the LISTSERV Maestro components always use the local names to communicate. When setting host names in INI files (or during the setup), use the names that are locally valid (which can also be externally valid names, provided the names work for local access, too). Whenever the local names are different from the external names (usually because some sort of proxying or forwarding is involved), the administrator needs to perform some additional configuration steps to make LISTSERV Maestro aware of the differences.



24.3.1 Configuring LISTSERV Maestro Components with Server Name Aliases or Proxies

The following list describes what to do if the host name, protocol, or HTTP-port used by external clients differs from the default values:

- **External host name, protocol, or port of the Maestro User Interface:**

The default host name is the value of the “HostName” setting in the `lui.ini`, or, if not present, the canonical host name of the server running the Maestro User Interface component. The default protocol is HTTP. The default port is “80” for HTTP and “443” for HTTPS.

If any of these values as used by external clients is different than these defaults, then you have to edit the **LUI Access URL for Users** setting in the Administration Hub to reflect this.

This may affect the **Default LUI Access URL for Users** under **Global Settings > Maestro User Interface > Default URL Settings** and/or the individual **LUI Access URL for Users** setting on group or user level.

You may also have to change the **LUI Access URL for Admin** setting (on the default level) and/or the **Subscriber Access URL** setting (on a group or user level), if these settings do not inherit the **LUI Access URL for Users** setting.

- **External host name, protocol, or port of the Administration Hub:**

The default host name is the value of the “HostName” setting in the hub.ini, or, if not present, the canonical host name of the server running the Administration Hub component. The default protocol is HTTP. The default port is “80” for HTTP and “443” for HTTPS.

If any of these values as used by external clients is different than these defaults, then you have to edit the **HUB Access URL for Users** setting in the Administration Hub to reflect this.

This may affect the **Default HUB Access URL for Users** under **Global Settings > Maestro User Interface > Default URL Settings** and/or the individual **HUB Access URL for Users** setting on group or user level.

You may also have to change the **HUB Access URL for Admin** setting (on default level), if it does not inherit the **HUB Access URL for Users** setting.

- **External host name, protocol, or port of Maestro Tracker:**

The default host name is the value of the “HostName” setting in the tracker.ini, or, if not present, the canonical host name of the server running the Maestro Tracker component. The default protocol is HTTP. The default port is “80” for HTTP and “443” for HTTPS.

If any of these values as used by external clients is different than these defaults, then you have to edit the **Tracking URL** setting in the Administration Hub to reflect this.

This may affect the **Default Tracking URL** under **Global Settings > Maestro User Interface > Default URL Settings** and/or the individual **Tracking URL** setting on group or user level.

See Section 5.3 [Setting the Default Tracking URL](#) for more details.

Using the two examples in the previous section (Section 24.3), you would have to make the following administration settings (assuming that you want to configure this as application defaults in the HUB):

- For **Example 1**, we would have to supply the following URL Settings:

Default LUI Access URL for Users: `http://maestro.sample.com:9001`

Default HUB Access URL for Users: `http://maestro.sample.com:9002`

Default Tracking URL: `http://maestro.sample.com:9003`

- For **Example 2**, we would have to supply the following URL Settings:

Default LUI Access URL for Users: `http://lui.sample.com`

Default HUB Access URL for Users: `http://hub.sample.com`

Default Tracking URL: `http://trk.sample.com`

(And then, of course, you would also have to configure the proxy accordingly so that it indeed does the transparent forwarding of the requests as described above – but this depends on the proxy software used and is not part of the LISTSERV Maestro setup.)



Note: If you change any of the access URLs in the HUB as described above (to allow for a different host name, protocol or port), and you make a mistake so that now the wrong access URL is submitted, then this may have the effect that users may no longer be able to login to LUI. This may even affect the admin login to LUI and HUB, in which case the admin can not even login normally to correct the mistake. See Section 31 [Emergency Admin HUB Access](#) for details on how to login into the HUB so you can correct the mistake.

Section 25 LISTSERV Maestro in Evaluation Mode

In order to function in normal mode, LISTSERV Maestro Lite needs to be connected to a fully licensed instance of LISTSERV. LISTSERV Maestro Lite connects to the configured LISTSERV instance to check if there is a MAESTRO scope and a suitably recent maintenance license in the LISTSERV license key (LAK). A “suitably recent” maintenance key is one that expires after the release date of the given LISTSERV Maestro Lite version. If not, then LISTSERV Maestro Lite will run in evaluation mode for all users and groups that are configured to use this LISTSERV instance.

In evaluation mode, actual delivery of a job is only simulated. When the scheduled send time of an authorized job has been reached, the job is transferred into the “delivered” state without actually sending any messages to any recipients. Operating in this fashion, a user can evaluate all aspects of LISTSERV Maestro Lite, including job definition, authorization, and viewing delivered jobs, without actually being able to use LISTSERV Maestro Lite for real deliveries.

Test delivery, which is a workflow step that precedes the authorization step, is possible even in evaluation mode. However, with the restriction that an evaluation message is added to the top of all messages that are sent during test delivery with LISTSERV Maestro Lite in evaluation mode.

In addition, for test delivery in evaluation mode to work, the following condition must be fulfilled:

- Either LISTSERV’s SMTP listener or a SMTP server must be installed on the **same server** on which the LISTSERV is running that is used for delivery, and this SMTP listener/SMTP server must be listening for SMTP requests on the standard SMTP port **25**.

Section 26 Adding Content to the Tomcat Server

26.1 Adding the LISTSERV Web Interface to the Tomcat Server

The Tomcat server of LISTSERV Maestro can also be used to serve the LISTSERV Web Interface (WA) when LISTSERV is installed on the same server as LISTSERV Maestro. This eliminates the need for an extra web-server for serving the web interface pages of LISTSERV.



Note: The following procedure is only meant for the situation where a single LISTSERV instance is present on the LISTSERV Maestro server, and Tomcat is supposed to serve the WA of this instance. If there are several ListPlex LISTSERV nodes on the server, see Section 26.4 [ListPlex and the Tomcat Server](#).

If LISTSERV Maestro and LISTSERV were installed together via the **Express Setup** option of the LISTSERV Maestro Setup-Suite, then this integration of the LISTSERV Web Interface into the Tomcat server has already been performed by the setup procedure.

If LISTSERV Maestro and LISTSERV are installed independently (but on the same server), then it is still possible to use LISTSERV Maestro's Tomcat server to also serve the LISTSERV Web Interface pages. The necessary folders are already in place in the Tomcat installation folder; therefore, the only additional configuration necessary is to tell LISTSERV about this.

This can be done either with the LISTSERV configuration tool or by manually editing the `site.cfg` file:

- Using the LISTSERV configuration tool:

Open the configuration tool, then click on **Advanced Configuration > Web Archives...**, and then enter the following information:

- In the **Directory from which the web server is authorized to run scripts** field, enter: `[install folder]\webapps\scripts\WEB-INF\cgi`
- In the **URL to use to access the web archive script in the above directory** field, enter: `/scripts/wa.exe`
- In the **Directory in which LISTSERV should place the files it creates and uses for the Web archive interface** field, enter:

`install folder]\webapps\archives`

(Replace "[install folder]" with the LISTSERV Maestro installation folder on your system.)

Finally, confirm your settings by clicking **[OK]**. Restart LISTSERV to make the changes effective.

- Editing the `SITE.CFG` file:

Add (or edit) the following entries:

```
WWW_ARCHIVE_CGI=/scripts/wa.exe
```

```
WWW_ARCHIVE_DIR=[install folder]\webapps\archive
```

```
SITE_CONFIG_CGI_DIR=[install folder]\webapps\scripts\WEB-INF\cgi
```

(Replace “[install folder]” with the LISTSERV Maestro installation folder on your system.)

Finally, save the `SITE.CFG` file with the new settings. Restart LISTSERV to make the changes effective.

- After the restart, LISTSERV should have copied the `wa.exe` file into the folder `[install folder]/webapps/scripts/WEB-INF/cgi`.

Verify that the account which is used to run LISTSERV Maestro (i.e. the account which is used to run the "L-Soft Tomcat" service) has execute rights on this `wa.exe` file. If not, add the execute right.

26.2 Adding Custom Content to the Tomcat Server

In LISTSERV Maestro, the various parts that are served by Tomcat are called “contexts”. Each context is an entity of its own inside of the Tomcat server. Each context has a name, which is also part of the URL that you use to access the context. More precisely, the context’s name is the part which appears right after the server name. For example, for LUI, the context name is `lui`, so the URL is `http://yourhost/lui`. Other LISTSERV Maestro contexts are `hub`, `trk`, `list`, `archives`, and `scripts`, with the respective URLs (`http://yourhost/hub`, `http://yourhost/list`, etc).

If you enter a context’s access URL (as above), what you actually get is the default page for that context (usually a page called `index.html`, `index.jsp`, `default.htm`, etc.). Therefore, if you type `http://yourhost/lui`, then what you actually get is the default page for the `lui` context (`http://yourhost/lui/index.jsp`).

To add content of your own (e.g. HTML pages, images, downloadable files, etc.) to the Tomcat installation of LISTSERV Maestro, you simply create a new context and put your files into that context. The files are then accessible using the URLs in that context.

Out of the box, Maestro does not support content or pages that are not part of a context. However, it is possible to support such content or pages once some additional configuration steps are taken (see Section 26.2.2 [Defining the Default Context](#)).

26.2.1 Adding Content as a New Context

To create a new context with your own content, the first question that you have to decide is the following:

“Does the server where Tomcat is running have several different host names, and if yes, do you want your own content to show up for all of these host names or not?”

Or in different words: Depending on which host name is used in the access URL when the user tries to access your content, is the content to show up for all host names in the URL, or only for specific ones?

The way that you proceed depends on how you answer this question:

- If the server has only a single host name anyway, or if you want your content to show up for all host names:

Create a new folder inside of the “webapps” folder of LISTSERV Maestro, like this:

```
[install_folder]/webapps/CONTEXT/
```

where you replace “CONTEXT” with the name of your context, for example:

```
[install_folder]/webapps/sample/
```

Then proceed as described further below.

- If the server has several host names, and you want your content to show up not for all of them, but only for one (or several) specific host name(s):

First you need to decide which one of the desired host names shall be the “main host name”. All others will be aliases. If you have only one desired host name, then this will be the “main host name” and there are no aliases.

Then create a new folder like this one:

```
[install_folder]/webapps-MAIN_HOST_NAME/CONTEXT/
```

where you replace “MAIN_HOST_NAME” with the “main host name” and “CONTEXT” with the name of your context, for example:

```
[install_folder]/webapps-host.domain.com/sample/
```

Now, add an entry like the following to the tomcat.ini file:

```
AdditionalHost.N=MAIN_HOST_NAME,ALIASLIST
```

where you replace “MAIN_HOST_NAME” with the “main host name” and “ALIASLIST” with a comma-separated list of all aliases (or leave out the “,ALIASLIST” part if there are no aliases). Also you need to replace the “N” with a unique number, i.e. there must not be two “AdditionalHost.N” entries with the same “N”. For example:

```
AdditionalHost.0=host.domain.com,alias.domain.com,alias.domain.org
```

```
AdditionalHost.1=host-without-alias.domain.com
```



Note: You can have several such “AdditionalHost” entries in the tomcat.ini file (as shown in the example above). Each of these entries defines one additional host, with a “main host name” and optionally a list of aliases for this host. The host names used by these entries must be unique, i.e. you must not use the same host name in two different “AdditionalHost” entries (neither as a “main host name” nor as an alias).

Then proceed as described below.

As your next step, copy the following folder (and the files in it) from the “archives” context to your own freshly created context (the archives context is automatically installed with each Maestro Tomcat). Copy the following folder:

```
[install_folder]/webapps/archives/WEB-INF/
```

So that at the end you have something like this:

```
[install_folder]/webapps/CONTEXT/WEB-INF/
```

or (in case you have a specific “main host name”):

```
[install_folder]/webapps-MAIN_HOST_NAME/CONTEXT/WEB-INF/
```

In the “WEB-INF” folder that you just copied, edit the file “web.xml” and look for the entry that says “<param-value>archives</param-value>”. Change the text “archives” as follows:

- If your context is in the default “webapps” folder, then change the text to your context name, like this: <param-value>CONTEXT</param-value>

- If your context is in a specific “webapps-MAIN_HOST_NAME” folder, then change the text to the “main host name” plus the context name, separated by a dash “-”, like this: `<param-value>MAIN_HOST_NAME-CONTEXT</param-value>`

Now you can put whatever files you want into your “CONTEXT” folder (you can also create subfolders). Usually you may want to include a start page like “index.html” or similar, but you can also have other pages, even in subfolders. Also image files or other downloadable content if you want.

Restart Maestro to make it aware of the new context.

The files in the new context will then be accessed for example as follows:

<code>http://HOST/CONTEXT</code>	access to the start page (if one was supplied)
<code>http://HOST/CONTEXT/</code>	also access to the start page (if one was supplied)
<code>http://HOST/CONTEXT/page.html</code>	access to page <code>page.html</code>
<code>http://HOST/CONTEXT/sub/other.html</code>	access to page <code>other.html</code> in the sub subfolder

Where of course you have to replace “CONTEXT” with your own context name and “HOST” with the correct host name for this context:

- If your context is in the default “webapps” folder, then you can use any host name that is assigned for this server to access this context, except for any host names which are used as a “main host name” or as an alias host name, as described above (if any).
- If your context is in a specific “webapps-MAIN_HOST_NAME” folder, then you can only use this “main host name” and its aliases to access this context. Also, you can not use this “main host name” or any of its aliases to access any contexts in the default “webapps” folder (including the default Maestro contexts) or any contexts in other “webapps-DIFFERENT_MAIN_HOST_NAME” folders (if there are any).



Warnings: In the default “webapps” folder, **do not** create custom contexts with one of the reserved names used by LISTSERV Maestro, i.e. do not call your context `lui`, `hub`, `trk`, `list`, `archives`, or `scripts`.

Do not put files into the `WEB-INF` folder, as they would not be accessible via a URL.

When creating subfolders in your context, **do not** create a folder called “`META-INF`”, as this is a reserved name

26.2.2 Defining the Default Context

The default context is the context that is used if the user types the access URL without any context name; for example, only “`http://yourhost`” or “`http://yourhost/`” or “`http://yourhost/somepage.html`”.

This is not supported out of the box by Maestro but must be configured first.

To define a default context, create a new context (with a new subfolder either in “webapps” or in “webapps-MAIN_HOST_NAME”) as described above.

In addition, before restarting Maestro, include the following entry in the `tomcat.ini`:

- If your context is in the default “webapps” folder:

```
DefaultContext=CONTEXT
```

- If your context is in a specific “webapps-MAIN_HOST_NAME” folder:

```
DefaultContext.MAIN_HOST_NAME=CONTEXT
```

where you replace “CONTEXT” with the name of your context and, if applicable, “MAIN_HOST_NAME” with the corresponding “main host name”.

For example for a context called “sample” you would have

```
DefaultContext=sample
```

- or -

```
DefaultContext.host.domain.com=sample
```



Warning: Do not define any of the reserved LISTSERV Maestro contexts as the default context (lui, hub, trk, list, archives or scripts)!

With this entry in the `tomcat.ini`, the specified default context can now be accessed in two ways:

- As before, you can still specify the context directly in the URL:

<code>http://HOST/CONTEXT</code>	access to the start page (if one was supplied)
<code>http://HOST/CONTEXT/</code>	also access to the start page (if one was supplied)
<code>http://HOST/CONTEXT/page.html</code>	access to page <code>page.html</code>
<code>http://HOST/CONTEXT/sub/other.html</code>	access to page <code>other.html</code> in the sub subfolder

- In addition, you can leave out the context in the URL but will still see the same pages (but this works only for the default context, of course):

<code>http://HOST</code>	access to the start page of the default context (if one was supplied)
<code>http://HOST/</code>	also access to the start page of the default context (if one was supplied)
<code>http://HOST/page.html</code>	access to page <code>page.html</code> in the root folder of the default context
<code>http://HOST/sub/other.html</code>	access to page <code>other.html</code> in the sub subfolder of the default context

26.3 Enabling Access Logging for Added Content

If you add your own content to be served by the Tomcat server, then you might also want to enable the usual web server logging for this content.

With enabled access logging, Tomcat will create log files in the same format as those created by standard web servers. These logs can later be analyzed by standard log analysis tools to track page hit counts, user session activity, and so on. The log files are created separately (under different names) for all contexts for which access logging is enabled and are rolled over nightly at midnight. All access log files are created in the following folder:

```
[install_folder]/logs
```

26.3.1 Enabling Access Logging for WA

To enable access logging for the WA component itself, edit the following file:

```
[install_folder]/webapps/scripts/META-INF/context.xml
```

and/or to enable access logging for the WA archives, edit the following file:

```
[install_folder]/webapps/archives/META-INF/context.xml
```

In this file, add a “<Valve>” tag similar to the following, just before the “</Context>” closing tag, so that the resulting “context.xml” file looks similar to this (the part that you should add is marked with bold):

```
<Context caseSensitive="false">
  <Manager className="org.apache.catalina.session.StandardManager" pathname="" />
  <Valve prefix="YOURNAME_access_log."
    className="org.apache.catalina.valves.FastCommonAccessLogValve"
    directory="logs" suffix=".log" pattern="common" resolveHosts="false" />
</Context>
```

In this tag, replace “YOURNAME” with a name that uniquely identifies the context for which you are enabling the logging, for example “WA” or “Archives” (this will become part of the log file name).

Then, restart LISTSERV Maestro.



Note: The above procedure applies to a normal LISTSERV installation only, with a single LISTSERV instance, whose WA is served by Tomcat. In case you are dealing with an installation, with several LISTSERV instances where Tomcat serves the WA applications of all of these instances (see Section 26.4 [ListPlex and the Tomcat Server](#)), then the procedure is slightly different.

In this case, the locations for the two “context.xml” files are slightly different. Instead of editing the files in the locations quoted above, edit the following files:

(1) To enable access logging for the WA component of a ListPlex node with a certain ListPlex-FQDN, edit the following file: [install_folder]/webapps-LSVNODE_FQDN/scripts/META-INF/context.xml

(2) And/or, to enable access logging for the WA archives of a ListPlex node with a certain ListPlex-FQDN, edit the following file: `[install_folder]/webapps-LSVNODE_FQDN/archives/META-INF/context.xml`

Also, when deciding on a substitution for “YOURNAME” (see above), you should include the FQDN of the LISTSERV node you are configuring (for example “LSVNODE_FQDN-WA” or “LSVNODE_FQDN-Archives”), so that the access logs of the WA applications of the various LISTSERV nodes will get different file names.

26.3.2 Enabling Access Logging for Custom Content

To enable access logging for a given custom content that you have added to Tomcat, use the following procedure:

- Copy the following folder (and the files in it) from the “archives” context to your own context (the archives context is automatically installed with each Maestro Tomcat). Copy this folder:

```
[install_folder]/webapps/archives/META-INF/
```

So that at the end you have something like this:

```
[install_folder]/webapps/CONTEXT/META-INF/
```

or (in case you have a specific “main host name”):

```
[install_folder]/webapps-MAIN_HOST_NAME/CONTEXT/META-INF
```

(where “CONTEXT” stands for the name of your context).

- In the “META-INF” folder that you just copied, edit the file “context.xml” and look for the “<Valve>” tag. If one already exists, edit the value of its “prefix” attribute. If no such “<Valve>” tag exists, add a new one similar to the following, just before the closing “</Context>” tag. The resulting “context.xml” file should then look similar to this (the part that you are supposed to add is marked with bold):

```
<Context caseSensitive="false">
  <Manager className="org.apache.catalina.session.StandardManager" pathname="" />
  <Valve prefix="CONTEXT_access_log."
    className="org.apache.catalina.valves.FastCommonAccessLogValve"
    directory="logs" suffix=".log" pattern="common" resolveHosts="false" />
</Context>
```

- In this tag, replace “CONTEXT” with the name of your context, or in case that you have a specific “main host name”, replace it with the “main host name” and the context name, separated by a dash “-” (this prefix will become part of the log file name).

Example for the context “sample” in the default folder “webapps”:

```
<Context caseSensitive="false">
  <Manager className="org.apache.catalina.session.StandardManager" pathname="" />
  <Valve prefix="sample_access_log."
    className="org.apache.catalina.valves.FastCommonAccessLogValve"
    directory="logs" suffix=".log" pattern="common" resolveHosts="false" />
</Context>
```

Example for the context “sample” in the specific folder “webapps host.domain.com”:

```
<Context caseSensitive="false">
  <Manager className="org.apache.catalina.session.StandardManager" pathname="" />
  <Valve prefix="host.domain.com-sample_access_log."
    className="org.apache.catalina.valves.FastCommonAccessLogValve"
    directory="logs" suffix=".log" pattern="common" resolveHosts="false" />
</Context>
```

- Finally, restart LISTSERV Maestro.

26.4 ListPlex and the Tomcat Server



Important: Setting up a server that has LISTSERV Maestro and multiple ListPlex nodes is a complex operation. For best results, we recommend that you utilize L-Soft’s consulting and installation services.

As described in Section 26.1 [Adding the LISTSERV Web Interface to the Tomcat Server](#), it is possible to use LISTSERV Maestro’s Tomcat server to serve the LISTSERV Web Interface (WA) for a LISTSERV instance that is installed on the same server as LISTSERV Maestro.

This is also possible if on the LISTSERV Maestro server there is not only a single LISTSERV instance, but a whole array of ListPlex nodes. The following describes how to set up the Tomcat server in this situation, to serve the WA for all nodes on the server.



Notes: The following procedure is only meant for the situation where Tomcat is supposed to serve several WAs at once. For the normal case with a single non-ListPlex LISTSERV installation and its WA, see Section 26.1 [Adding the LISTSERV Web Interface to the Tomcat Server](#).

The procedure described here assumes that you are starting with a clean server, on which neither LISTSERV Maestro or any LISTSERV instance is installed yet. For example, a server where you freshly set up a LISTSERV Maestro installation plus a number of ListPlex nodes. However, the procedure can easily be adapted to situations where LISTSERV Maestro is already installed and addition ListPlex nodes will be added, or the other way round, where some ListPlex nodes are already installed and LISTSERV Maestro is added and is supposed to serve the WAs for these nodes.

The ability to run multiple ListPlex nodes on the same server is only available on Windows; therefore, the information in this section only applies to those using Windows.

1. Preparation

For this procedure to work, it is necessary that for the server in question you have several fully qualified domain names (FQDNs) as well as several IP-addresses.

One set of FQDNs and IP-addresses is required for each ListPlex node:

- For each ListPlex node, you need one separate FQDN whose DNS entry points to a separate IP-address that is assigned to this server (i.e. all ListPlex nodes must have dedicated FQDNs and IP-addresses). The same IP-address should not be used by any other FQDN. Below, we call these FQDNs the “ListPlex-FQDNs”.

One of the reasons for the dedicated ListPlex-FQDNs is, that this makes it possible for each ListPlex customer to access his personal ListPlex node by using his personal ListPlex-FQDN.

The question now is, do you want to offer the same behavior when accessing the Maestro User Interface? Depending on how you answer this, you need slightly different setups:

- If all users are supposed to access the Maestro User Interface using the same shared FQDN, then you need one additional FQDN.

Below, we call this one additional FQDN the “shared-LMA-FQDN”.

The shared-LMA-FQDN must not reuse any of the ListPlex-FQDNs, but must be a separated FQDN. However, it does not have to point to a separate IP-address. It can point to one of the already existing addresses (for example also share it with one of the ListPlex-FQDNs of above).

- If you want all users to access the Maestro User Interface with the same personal FQDN that they also use to access the ListPlex nodes, then you do not need any further FQDNs. Each user will then access LISTSERV Maestro with the same ListPlex-FQDN as he uses to access the WA.

The dedicated-LMA-FQDNs must not reuse any of the ListPlex-FQDNs, but each must be a separate FQDN. However, these dedicated-LMA-FQDNs do not have to use separate IP-addresses, but can for example share the dedicated IP-address of the corresponding ListPlex-FQDN of above. The best practice for this is, to define a dedicated IP-address per ListPlex node (as described above) and then define the ListPlex-FQDN and the dedicated-LMA-FQDN for this customer to both point to this dedicated IP-address.

2. Install the ListPlex Nodes

First, install the ListPlex nodes as you usually would. I.e. follow the usual procedure about how to set up several ListPlex nodes on one server. When installing the ListPlex nodes, do not set them up for using the WWW Interface (WA) at this time (this will come later, as explained below). That is, install all nodes as if no WA access was required.

When installing the nodes, use the dedicated ListPlex-FQDN with the dedicated IP-address for each node, as is customary when setting up ListPlex nodes (see above).



Note: Procedures for installing multiple ListPlex nodes on a single server are available from L-Soft Support.

3. Install LISTSERV Maestro

Install a fresh instance of LISTSERV Maestro on the server. Use the following options and choices in the install wizard:

- On the first Setup Type screen, select **Custom Setup**.
- On the LISTSERV Maestro Package: Individual Products Setup screen, click on **LISTSERV Maestro**.
- On the second Setup Type screen, select **Custom Setup** again.
- On the Existing LISTSERV Installation Found screen, select the **No** option.
- On the Default LISTSERV Settings screen, leave all entries empty for now. The LISTSERV connection settings will later be supplied in the HUB.

4. Configure Tomcat to serve the WA application for all ListPlex Nodes

The following procedure must be repeated once for each of the ListPlex nodes:



Note: In the following, certain placeholders are used which need to be replaced as follows:

`ListPlex_FQDN` – Replace with the ListPlex-FQDN of the node that you are currently configuring.

`[install_folder]` – The installation folder for LISTSERV Maestro.

- Create a folder with the following name:

```
[install_folder]/webapps-ListPlex_FQDN
```

- Copy the following two folders (and the subfolders and files in them) into the “webapps-ListPlex_FQDN” you created above, so that this folder then contains an “archives” and a “scripts” subfolder. Copy these folders:

```
[install_folder]/webapps/archives
```

```
[install_folder]/webapps/scripts
```

to create these folders:

```
[install_folder]/webapps-ListPlex_FQDN/archives
```

```
[install_folder]/webapps-ListPlex_FQDN/scripts
```

- Edit the following file:

```
[install_folder]/webapps-ListPlex_FQDN/archives/WEB-INF/web.xml
```

Look for the entry that says “<param-value>archives</param-value>”. Edit the value of this entry so that the entry instead looks like this:

```
<param-value>ListPlex_FQDN-archives</param-value>
```


- Edit the following file:

```
[install_folder]/webapps-ListPlex_FQDN/scripts/WEB-INF/web.xml
```

Look for the entry that says "<param-value>scripts</param-value>". Edit the value of this entry so that the entry instead looks like this:

```
<param-value>ListPlex_FQDN-scripts</param-value>
```



- Create four directory junctions ([NTFS junction points](#)) in the following folder:

```
[install_folder]/webapps-LISTPLEX_FQDN/
```

The first junction must be called "lui" and must point to:

```
[install_folder]/webapps/lui/
```

The second junction must be called "hub" and must point to:

```
[install_folder]/webapps/hub/
```

The third junction must be called "trk" and must point to:

```
[install_folder]/webapps/trk/
```

The fourth junction must be called "list" and must point to:

```
[install_folder]/webapps/list/
```

To create these junctions, do so here:

On Windows 2000, Windows XP, Windows Server 2003:

- You need the tool linkd.exe from the Windows Resource Kit (for example the "Windows Server 2003 Resource Kit", can be downloaded from Microsoft).
- Open a command line window and use the cd command to change into the following directory: [install_folder]/webapps-LISTPLEX_FQDN/
- Then type (submit with ENTER):
linkd lui "[install_folder]/webapps/lui"
- This creates the junction called "lui". Repeat to also create the "hub", "trk" and "list" junctions.

On Windows Vista, Windows 7, Windows Server 2008:

- The necessary tool mklink.exe is already part of the operating system.
- Open a command line window and use the cd command to change into the following directory: [install_folder]/webapps-LISTPLEX_FQDN/
- Then type (submit with ENTER):
mklink /J lui "[install_folder]/webapps/lui"
- This creates the junction called "lui". Repeat to also create the "hub", "trk" and "list" junctions.

Junctions behave more or less like normal folders. But any changes you make in one of these folders will also affect the original folder in "[install_folder]/webapps" and vice versa.

- Edit the following file:

```
[install_folder]/conf/tomcat.ini
```

Add an entry like the following:

```
AdditionalHost.N=ListPlex_FQDN
```

where you replace the “N” with a unique number (for example “0”, or “1” or similar), in such a fashion, that if there are several “AdditionalHost.N” entries in the `tomcat.ini` file (as is the case as soon as you configure the second ListPlex node), then each entry must use a different value for “N”.

- Copy the file `wa.exe` from the “MAIN” folder of the ListPlex node to the following folder:

```
[install_folder]/webapps-LISTPLEX_FQDN/scripts/WEB-INF/cgi
```



- Verify that the account used to run LISTSERV Maestro (i.e. the account used to run the "L-Soft Tomcat" service) has execute rights on this `wa.exe` file. If not, add the execute right.

- Skip this step if you are currently configuring the main LISTSERV instance. Only do this step if you are currently configuring one of the secondary LISTSERV instances:

Rename the `wa.exe` file that you just copied into the “cgi” folder from “`wa.exe`” to “`wa-INSTANCE_NAME.exe`”, where “INSTANCE_NAME” must be replaced with the value of the “INSTANCE” variable in the `SITE.CFG` of the ListPlex node that you are currently configuring. For example “`wa-hq.exe`”.

- Edit the `SITE.CFG` of the ListPlex node that you are currently configuring and add the following entries (linebreak added for readability only):

```
WWW_ARCHIVE_CGI=/scripts/WA_EXE_NAME
```

```
WWW_ARCHIVE_DIR=[install_folder]\webappsListPlex_FQDN\archives
```

```
SITE_CONFIG_CGI_DIR=[install_folder]\webapps-ListPlex_FQDN\scripts\WEB-INF\cgi
```

where you replace “WA_EXE_NAME” with the name that you gave the WA executable in the previous step (make sure to check which name the WA executable for this ListPlex node has in the `[install_folder]\webapps-ListPlex_FQDN\scripts\WEB-INF\cgi` folder).



Important: In the above entries in the `SITE.CFG`, make sure not to specify a path which contains spaces. This means, that if the path of “[install folder]” contains folders with spaces in their names (which is, for example, the case for LISTSERV Maestro’s default install folder “`C:\Program Files\L Soft\Application Server\`”), then you must use the space-free 8.3 variants of these folder names, for example “`C:\Progra~1\L-Soft\Applic~1\`”.

- Repeat the above steps for the next ListPlex node, until all nodes have been configured in the same fashion. If you later add a new node, repeat the same steps for this node too.

5. Restart all LISTSERV Instances and LISTSERV Maestro.



6. Configure the LUI Access URLs

The LUI access URLs that are configured in the Administration Hub must now also be modified to include the FQDNs used above.

- If a shared-LMA-FQDN is used, then you should simply configure the **Default LUI Access URL** and let all accounts/group inherit this default setting (since they are all using the shared-LMA-FQDN). Edit the **Default LUI Access URL** to use this shared-LMA-FQDN. To edit, select **Global Settings > Maestro User Interface > Default URL Settings**. All accounts and groups should then simply inherit this setting.
- If the ListPlex-FQDNs are used even to access LISTSERV Maestro, then each account/group has its own FQDN, so each has a different LUI access URL too. Because of this, you need to edit the **LUI Access URL** of each account/group to use the corresponding ListPlex-FQDN. To access, select the user account or group, then **User Account** or **Group > URL Settings**. Edit the **LUI Access URL for Users** or **LUI Access URL for Group**.

LISTSERV Maestro and the LISTSERV Web Interface for each of the nodes will now be available with a web browser via the following URLs:

To access LISTSERV Maestro:

- If a shared-LMA-FQDN is used, then for all users, the Maestro User Interface is accessible via this shared name, with the following URL:

```
http://SHARED_LMA_FQDN/lui
```

- If the ListPlex-FQDNs are used even to access LISTSERV Maestro, then for each user, the Maestro User Interface is accessible via his personal dedicated FQDN, with the following URL:

```
http://LISTPLEX_FQDN/lui
```

To access the LISTSERV interface (WA):

- For each user, WA is accessible via its personal LISTSERV-FQDN only, with the following URLs:

```
http://LISTSERV_FQDN/archives
```

```
http://LISTSERV_FQDN/scripts/wa.exe
```

With this setup, you now have the following behavior/restrictions:

- To access the Web Interface of a certain ListPlex node, you **must** use the ListPlex-FQDN of that node. Using one of the other ListPlex-FQDNs or the shared-LMA-FQDN (if any) will not work.
- When using a shared-LMA-FQDN, you **must** use this FQDN to access the Maestro User Interface. Using one of the ListPlex-FQDNs will not work.

- When using the ListPlex-FQDNs to access LISTSERV Maestro, then to access it with a certain user account, you **must** use the ListPlex-FQDN associated with this account (or its group). Using one of the other ListPlex-FQDNs will not work.

Section 27 Using International Character Sets

Each email job that is created in LISTSERV Maestro has a character set (charset) associated with its content. This charset is used to encode the content for sending. When a job is first created as a new job (not as a copy of an existing job), the job is initially created with the default charset. LISTSERV Maestro defaults to the ISO-8859-1 (Latin 1) character set for encoding email messages unless the administrator has defined a different default setting.

27.1 Defining the Default Mail Charset

To define the default charset, edit the following in the Maestro User Interface INI-file:

```
[maestro_install_folder]/lui/lui.ini
```

Edit or add the key `DefaultMailCharset` and set it to the name of one of the charsets supported by LISTSERV Maestro.

Table 27-1 Supported Charsets

Charset Name:	Description:
US-ASCII	US ASCII
ISO-8859-1	West European, Latin 1
ISO-8859-2	East European, Latin 2
ISO-8859-3	South European, Latin 3
ISO-8859-4	North European, Latin 4
ISO-8859-5	Cyrillic
ISO-8859-6	Arabic
ISO-8859-7	Greek
ISO-8859-8	Hebrew
ISO-8859-9	Turkish, Latin 5
ISO-8859-15	Similar as ISO-8859-1 but with Euro currency symbol
UTF-8	International Unicode, encoded in UTF-8 format
GB2312	Simplified Chinese (GB2312)
BIG5	Traditional Chinese (BIG5)
ISO-2022-JP	Japanese (ISO-2022-JP)
X-EUC-JP	Japanese (X-EUC-JP)
X-SJIS	Japanese (X-SJIS)

KSC_5601	Korean (KSC_5601)
EUC-KR	Korean (EUC-KR)
AUTO-NO-UTF-8	<p>LISTSERV Maestro will choose either US-ASCII or any of the ISO-8859 charsets (but <i>not</i> UTF-8), depending on the characters that are actually used in the content. If possible, ASCII is favored over any ISO-8859, so an ISO-8859 set is only chosen if ASCII is not able to display all characters in the content.</p> <p>Of the ISO-8859 sets, the one where the number of non-displayable characters is minimized is chosen. If two sets have an equal number of non-displayable characters, then lower ISO-8859 sets are favored over higher sets (for example, ISO-8859-1 over ISO-8859-2, over ISO-8859-3, and so on).</p>
AUTO-YES-UTF-8	<p>LISTSERV Maestro will choose either US-ASCII or any of the ISO-8859 or even UTF-8, depending on the characters that are actually used in the content. If possible, ASCII is favored over any ISO-8859 and the ISO-8859 sets are favored over UTF-8.</p> <p>The step to the next “higher” set is only made if the “lower” set is not able to display all characters in the content. If several ISO-8859 sets are able to display all characters, then lower ISO-8859 sets are favored over higher sets (for example ISO-8859-1 over ISO-8859-2, over ISO-8859-3, and so on.).</p>

The default charset is only initially assigned to the email job. Users may change the default charset on the Define Content screen.

If the administrator wants to prevent the users from changing the default charset (and force the users to always accept the default charset already set), another entry in the same INI-file needs to be edited:

- Edit or add the key `AllowCharsetChoice`. Set to `true`, allowing the users to change the charset of a job (to be able to assign different charsets to each job) or to `false` to disallow changing of the charset. The default if the key is not present in the INI-file is `true`.

27.2 Allowing or Disallowing Bi-Directional Character Sets

Of the ISO-8859 charset family, two charsets contain letters from languages that have a standard reading direction of right-to-left. These are the charsets ISO-8859-6 (Arabic) and ISO-8859-8 (Hebrew), both of which are supported by LISTSERV Maestro.

Actually, LISTSERV Maestro will not use the charsets with the names ISO-8859-6 and ISO-8859-8 when it recognizes Arabic or Hebrew characters, but will instead use the special bi-directional versions *ISO-8859-6-i* and *ISO-8859-8-i*. These charsets contain the same characters as their non-i-suffix counterparts, but the “-i” suffix tells the receiving mail client that the text should be displayed with right-to-left reading direction. Without the “-i” suffix in the charset name, many email clients would probably display the correct characters, but in the (for that language) incorrect left-to-right reading direction.

Even with the “-i” suffix, the recipient might need a special mail client version (or even a special mail client) that is prepared to display text with right-to-left reading direction

properly and is also able to properly display bi-directional text (text that mixes characters with left-to-right and characters with right-to-left reading direction, in the case of a Hebrew text that contains English names, for example). Some clients may only display the characters with the right direction, but still left-align each line of text, instead of the correct right-alignment (occurrences such as this are subject to the mail client itself, and are outside of the scope of LISTSERV Maestro).

It is possible, however, to disallow the charsets with the "-i" suffix and use the "normal" counterparts, ISO-8859-6 and ISO-8859-8 instead. To do so, edit the following file:

```
[maestro_install_folder]/lui/lui.ini
```

Edit or add the key `AllowISO-i-Mails=false` to disallow the bi-directional charsets. (If the key from the INI-file is removed, commented out, or set to `...=true`, then the bi-directional charsets will be allowed as the default).

This INI-file setting will affect all mail sent, with any user account. Please note that changing this setting requires a restart of the Maestro User Interface component to take effect.

Section 28 LISTSERV Maestro Standard Default Ports

The following tables list the standard default ports for the Administrative Hub, the User Interface, and the Maestro Tracker.

Ports Used by the Administrative Hub

Port Number	Function
80	HTTP access
1099	Internal communication with other LISTSERV Maestro Components
8007	Shut down of the application server

Ports Used by the Maestro User Interface

Port Number	Function
80	HTTP access
1099	Internal communication with other LISTSERV Maestro components
8007	Shut down of the application server
3306	Internal system database connection

Ports Used by Maestro Tracker

Port Number	Function
80	HTTP access
1099	Internal communication with other LISTSERV Maestro components
7000	Communications Port transfers tracking data to the Maestro User Interface
8007	Shut down of the application server

Section 29 Updating Maestro's HTML Upload Applet

The HTML upload applet that is part of LISTSERV Maestro is digitally signed with a digital certificate, which is a mechanism to authenticate the trustworthiness of the applet before allowing it to access the user's local disk (for loading the HTML file and inline binaries). The digital certificate that is used for this process is issued by a trustworthy signing authority and has a limited validity duration, i.e. it expires at a given date. (To date, L-Soft uses certificates issued by Thawte Consulting or VeriSign.)

After the certificate has expired, the upload applet will still continue to function but the user may be presented with a warning message that the certificate has expired.



Note: An upload applet with an expired certificate is just as secure as an upload applet with a valid certificate; therefore, if you are not bothered by the warning messages, you may continue to use the expired applet without any worries. However, the warning message about the expired certificate may irritate users and leave them uncertain about if it is safe to execute the applet. If you want to avoid this, then you need to replace the expired applet with a freshly signed version as described below.

The upload applet comes in two versions: One version is for the Sun Java-Plugin, which is used by many modern browsers (for example by Firefox, Safari and Mozilla, and often also by Internet Explorer); the other version is for the Microsoft Java Virtual Machine, which is used mostly by older Internet Explorer installations. The version of the applet that is being used depends entirely on the browser used to access LISTSERV Maestro, (i.e. at one and the same LISTSERV Maestro installation, both applet versions may be used). The two versions of the applet are identical in terms of their function and how they work, but the certificate messages they present to the user are different, as described in the following sections.

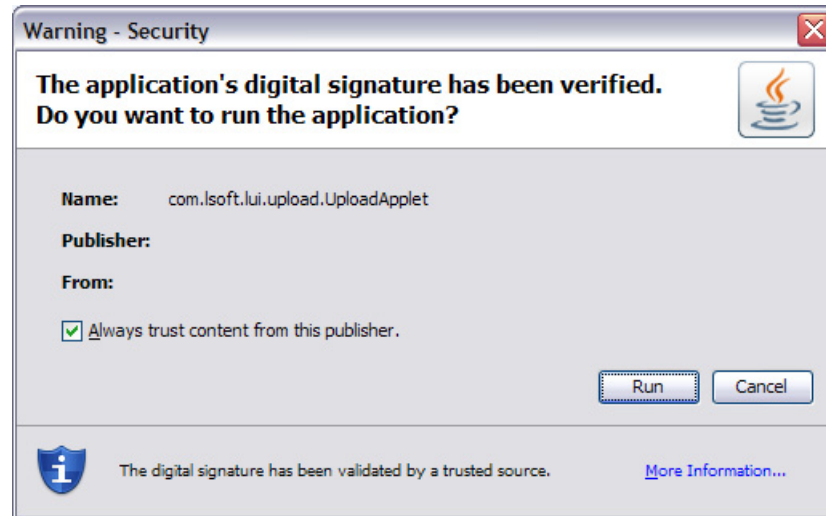
29.1 Sun Java-Plugin

If the certificate has expired, a message like the following will be displayed when the upload applet is started for the first time in a given browser session:



(If you click "Yes", then the upload applet will still function normally.)

For an applet with a certificate that is still valid, the message will look similar to this:



If you are seeing the former of the two messages, then the certificate of the applet in your LISTSERV Maestro has expired. If you want to avoid the warning message so that the second message is being displayed instead, you need to update the applet in your LISTSERV Maestro server with a version that has been signed with a fresh certificate that is still valid.

Please contact L-Soft support with information about the LISTSERV Maestro version you are using and you will be supplied with a freshly signed version of the applet (it might be necessary that you update to the newest LISTSERV Maestro version before you can apply the new applet). The applet comes in form of a file called "upload.jar", which you need to copy to the following location (overwriting the file of the same name which is already present):

[Maestro_install_folder]/webapps/lui/JSP

It is not necessary to restart LISTSERV Maestro after this replacement, but if you already had a browser open and were accessing LISTSERV Maestro, you may need to close and restart the browser to make it aware of the new applet.



Note: The Sun Java-Plugin sometimes caches the applet on the local computer. This may have the effect that on a given computer the old expired version of the applet is being used even after you have replaced the "upload.jar" file. If this is the case, then close your browser and clear the applet cache of the Sun Java-Plugin. For example, on Windows and with the "Java 6" Java-Plugin: Select the Windows **Start Menu**, click on **Control Panel**, click on the **Java** icon, click on the **General** tab, click the **[Settings]** button, click the **[Delete Files]** button, check the **Applications and Applets** option, and then click **[OK]**.

29.2 Microsoft Java Virtual Machine

The MS Java VM does not display an explicit warning message for an expired certificate; therefore, if this applet version is being used, users may not even notice that a certificate has expired. However, if the certificate details are examined closely, the fact that the applet is expired may still be noted.

If the applet is started for the first time in a given browser session, a message similar to the following is displayed:



This message looks exactly the same both for applets with a valid or an expired certificate. If the user now clicks **[Yes]**, then the applet will start normally without the user even noticing that the certificate may have expired.

If you click on **L-Soft Germany GmbH**, then the following window is displayed where the certificate validity can be seen (i.e. here a user may notice that the certificate is expired):



If you want to prevent a user from discovering that the applet has already expired, then you need to update the applet in your LISTSERV Maestro server with a version that has been signed with a fresh certificate that is still valid.

Please contact L-Soft support with information about the LISTSERV Maestro version you are using to receive a freshly signed version of the applet (you may need to update to the newest LISTSERV Maestro version before you can apply the new applet). The applet comes in form of a file called "upload.cab", which you need to copy to the following location (overwriting the file of the same name which is already present):

```
[Maestro_install_folder]/webapps/lui/JSP
```

It is not necessary to restart LISTSERV Maestro after this replacement, but if you already had Internet Explorer open and were accessing LISTSERV Maestro, then you may need to close and restart Internet Explorer to make it aware of the new applet.

Section 30 LISTSERV Maestro Lite

LISTSERV Maestro supports a special “lite-mode”. In this context, “lite” means that certain features are restricted (not available). The restricted features can be grouped into two categories:

- **User-feature:** Features that apply to the actual LUI users, i.e. such a restriction means that the LUI user has less options/features than in the full version. For example “No drop-ins” or “No personal tracking” are restrictions that fall into this category.
- **Admin-feature:** Features that apply to the actual LISTSERV Maestro installation and the LISTSERV Maestro administration, i.e. such a restriction means that the LISTSERV Maestro admin has less options/features than in the full version. For example “No external system database” or “No dedicated bounce server” are restrictions that fall into this category.

The lite-mode itself can be activated in two different ways:

- **LISTSERV Maestro in lite-mode:** This means that the whole LISTSERV Maestro installation is running in lite-mode. In this case, both the user-feature and admin-feature restrictions apply, and they apply for all users, as well as the admin. To activate this mode, add the following entry to the hub.ini file:

```
LiteMode=true
```



Note: Configuring LISTSERV Maestro to run in lite-mode as described above is a precondition for lite-LAKs to be accepted during LUI login. If LISTSERV Maestro is in lite-mode, then for a LUI login either a full-LAK or a lite-LAK will be accepted. However, in full-mode, only a full-LAK will be accepted. If only a lite-LAK is found while running in full-mode, then the user will be logged in with evaluation mode just as if no LAK was present at all.

- **Specific users in lite-mode:** This means that LISTSERV Maestro is running in full-mode, but that for certain users the lite-restrictions do still apply. In this case, only the user-feature restrictions apply, and only to specific users (not necessarily all users). The admin-feature restrictions do not apply. To activate this mode, do not add the “LiteMode=true” setting in the hub.ini (so that Maestro by itself is running in full-mode), but instead add the following entry to the hub.ini file:

```
LiteModeConfigurableForScopes=true
```

With this setting, a new configuration section becomes available in the Administration Hub (under the Maestro User Interface settings), which is normally hidden. It is called **Lite-Mode Restrictions** (on group or user level) or **Default Lite-Mode Restrictions** (on application default level).

In this section, you can define if the account (or group) to which the setting applies will have the lite-mode restrictions enforced or not. If you enforce the restrictions on

default application level, then this applies to all groups and accounts that do not specify a different setting on the individual level.

All accounts for which the lite-mode restrictions are enforced in this manner will have all the user-feature restrictions enforced. All other accounts/groups will be able to use the full feature set and the admin-feature restrictions will also not be enforced.



Note: As described above, LISTSERV Maestro must run in full-mode when using user specific lite-mode enforcement. This means, that even for users where lite-mode enforcement is switched on, a full-LAK must be available during LUI login, because in full-mode, only full-LAKs are accepted.

30.1 Current Lite-Mode Restrictions

The following sub-sections list the feature restrictions that are currently enforced by lite-mode, in their respective “user-feature” or “admin-feature” categories.

30.1.1 User-Feature Restrictions

User feature restrictions are restrictions that apply to users who log in at LUI. These restrictions apply both to normal LUI users, as well as to the admin if he logs in at LUI, unless noted otherwise.

30.1.1.1 Mail Jobs

Mail job restrictions:

Feature	Standard	Lite
Standard mail jobs	X	X
A/B-split jobs	X	
Job categories (only available in 4.0)	X	
Job-list filters	X	X
Forward-To-A-Friend	X	
Auto-repeat jobs	X	
Team collaboration settings on jobs	X	
X-headers in sender definition	X	
DKIM support	X	
Plain text content	X	X
HTML content with plain text alternative	X	X
HTML content with AOL-format alternative	X	
Conditional content parts for HTML content with alternative	X	
Job-archivation by user	X	

Feature	Standard	Lite
Job-archivation by admin	X	X
Auto-job-archivation settings by user	X	
Auto-job-archivation settings by admin	X	X
Multi-job-actions for user	X	
Multi-job-actions for admin	X	X
Authorization-due-by date	X	
Target Group Recipients	X	X
LISTSERV List Recipients	X	X
Recipients Based on Reaction to the Job	X	
CSV recipients	X	
"Select from database" recipients	X	
"Let LISTSERV select from database or LDAP directory" recipients	X	
Job journal	X	
Manual trigger for jobs with delivery schedule "Send when triggered"	X	X
External trigger for jobs with delivery schedule "Send when triggered"	X	
Support for "View-in-Browser" URL	X	

30.1.1.2 Tracking

Tracking restrictions:

Feature	Standard	Lite
Personal tracking	X	
Anonymous tracking	X	
Unique tracking	X	X
Blink tracking	X	X
"Open-up" tracking	X	X
"Track all links"	X	X
"Track no links"	X	X
Individual link tracking	X	
Link aliases when tracking	X	
Action tracking	X	
Permission based tracking	X	

30.1.1.3 Reporting

Reporting restrictions:

Feature	Standard	Lite
Job comparison reports	X	
Job statistics reports	X	
Recipients statistics reports	X	
Tracking reports (report administration by admin not restricted)	X	
Quick reports	X	X
Reports on the dashboard	X	

30.1.1.4 Recipient Warehouse

Recipient Warehouse restrictions:

Feature	Standard	Lite
Number of datasets	No limit	1
Number of hosted lists in dataset	No limit	No limit
Linked LISTSERV lists in dataset	X	X
Categories in datasets	X	
Team collaboration settings on datasets	X	
Target groups based on dataset	X	X
Target groups based on hosted list	X	X
Target groups based on standard LISTSERV list	X	X
CSV recipients target groups	X	
"Select from database" recipients target groups	X	
"Let LISTSERV select from database or LDAP directory" recipients target groups	X	
Conditions and parameters in HRL and LISTSERV list target groups	X	
Derived profile fields in datasets or hosted lists	X	
Direct import into dataset or hosted list from a CSV file	X	X
Direct import into dataset or hosted list from an external database	X	
Predefined member importers or recipient importers	X	
External download of dataset members or hosted list recipients	X	

Feature	Standard	Lite
Lookup tables	X	X
Lookup table subsets	X	
Lookup table secondary values	X	

30.1.1.5 Utility

Utility feature restrictions:

Feature	Standard	Lite
Freely created user defined content templates	X	
Template creation as copies of system templates (with less system templates to choose from) or by template upload (only templates downloaded from a lite-account or specially prepared templates can be uploaded)		X
Fully editable user defined content templates	X	
Templates can only be edited in minor ways: Name, subject, certain attributes. (Template administration by admin not restricted)		X
System-Drop-ins	X	X
User-Defined Drop-ins (drop-in administration by admin not restricted)	X	
Sender profiles (profile administration by admin not restricted)	X	
Multiple dashboards	X	

30.1.2 Admin-Feature Restrictions

Admin feature restrictions are restrictions that apply to the administration and the installation of LISTSERV Maestro as a whole.

Feature	Standard	Lite
Number of User Scopes (Scope: a single user or group, with any number of group members)	No limit	1
User Identities	X	
Option to revoke "create job" right from users	X	
Option to define alternative "designated job owner" for users	X	
Single server install (LUI, HUB, TRK all on one server)	X	X
Distributed install (LUI, HUB, TRK not all on one server)	X	
External system database	X	

Feature	Standard	Lite
Dedicated bounce server	X	
Daily/monthly mail limits	X	
System metrics reports and daily system metrics notification emails	X	

Section 31 Emergency Admin HUB Access



Login of the admin to the Administration Hub normally happens via the Maestro User Interface (LUI) login page. Under some circumstances, it may however happen that login to LUI is not possible so that the admin can not login either. This means that the admin also does not have access to the HUB. Now, if the reason for this problem is some sort of incorrect configuration in the HUB, for which the admin needs to login to the HUB to be able to fix the problem, then there is a catch-22 situation: The admin can not login to the HUB because of the problem, but to be able to fix it, he needs to login to the HUB.

For such a situation, there is an emergency login access directly to the HUB, via the HUB login page (instead of the LUI login page that is normally used). Normally, the HUB login page would not allow the admin to login but would instead redirect him to the LUI login page.

However, the HUB login page can be forced to allow a direct login, circumventing any problems with the LUI login. To force such a direct login, access the HUB login page with the following URL:

```
http://HOSTNAME/hub?loginOverride
```

(where you replace "HOSTNAME" accordingly).

Then submit the admin password to login. You will now be logged in to the HUB, however without the option of switching over to the LUI interface.

Section 32 Sharing a Server with IIS



If you want to install LISTSERV Maestro on a server where an IIS web server is already running, then most likely the default HTTP port 80 is already in use by this web server. Therefore, you need to configure LISTSERV Maestro to use a different HTTP port. For details on configuring LISTSERV Maestro to use a different HTTP Port, see Section 14.2.1 [Ports Used by LISTSERV Maestro](#).

The unfortunate side-effect of this is that users of LISTSERV Maestro will be unable to access Maestro with a nice "clean" URL, like `http://myhost.domain.com/loi`. Instead, they would have to include the non-standard HTTP port in the URL, like `http://myhost.domain.com:8080/loi`.

While this may be acceptable in many cases, in some cases it is not. In these cases, it is possible to "share" the standard port 80 with IIS or Apache.

Or more precisely: IIS acts as a reverse-proxy in such a way that the requests from the clients are made to IIS (on the default port), and IIS then dispatches the requests to LISTSERV Maestro as the backend server. LISTSERV Maestro's response is then passed back to IIS, which forwards it to the client. To the client this looks, as if it is communicating directly with LISTSERV Maestro on the standard port, when, in reality, it is communicating with IIS. However, this does not process the requests itself, but passes them on to LISTSERV Maestro and then passes the responses back to the client.

The following sections describe how to configure LISTSERV Maestro and IIS for such a scenario.



Important: The web site that is to be used for the proxying must not have any root folders with the names "loi", "hub", "list" or "trk", i.e. before configuring the web site as described below, make sure that the following URLs are **not** in use (and also not any sub-URLs that start with these URLs):

```
http://YOURSERVER/loi
http://YOURSERVER/hub
http://YOURSERVER/list
http://YOURSERVER/trk
```

32.1 Configuring LISTSERV Maestro

1. Edit the file `[install_folder]/conf/tomcat.ini` and add the following entry:

```
AJPConnector=IP_ADDRESS:PORT
```

where you replace `IP_ADDRESS` with one of the IP-addresses of the server (usually the main address) and `PORT` with a free port that is not used by any other application (for example, 8009).

Example:

```
AJPConnector=192.168.1.1:8009
```

2. Create a folder on the server, in a suitable location, with a suitable name. The account under which the web server is running must have “read and write” access to this folder. In the following steps, we will use `[connector_folder]` to depict this folder.
3. Create a text file with the following name:
`[connector_folder]/workers.properties`
4. Edit this file with a text editor and set its content to the following four lines:
`worker.list=maestro`
`worker.maestro.type=ajp13`
`worker.maestro.host=HOSTNAME`
`worker.maestro.port=PORT`
where you replace `HOSTNAME` and `PORT` according to the values that you used for the `AJPCoconnector` setting in the `tomcat.ini`. The `HOSTNAME` must be a host name that maps to the IP-address that you used for the `AJPCoconnector`, and the `PORT` must be the same port as in the `AJPCoconnector`.
5. Create a text file with the following name:
`[connector_folder]/uriworkermap.properties`
6. Edit this file with a text editor and set its content to the following four lines:
`/lui/*=maestro`
`/hub/*=maestro`
`/trk/*=maestro`
`/list/*=maestro`

32.2 Configuring IIS

1. Download the latest stable version of the ISAPI Redirect DLL:
For 32-bit Windows:
<http://apache.copahost.com//tomcat/tomcat-connectors/jk/binaries/win32/>
For 64-bit Windows:
<http://apache.copahost.com//tomcat/tomcat-connectors/jk/binaries/win64/>
You only need to download the actual DLL from the release folder. It has a name like `isapi_redirect-VERSION.dll`. At the time of writing, the latest stable release was 1.2.31, so the filename was `isapi_redirect-1.2.31.dll`.
Download this file and then rename it to remove the version number from the name, so that it then has the name `isapi_redirect.dll`.
Put this file into the `[connector_folder]`.
2. Create a text file with the following name:
`[connector_folder]/isapi_redirect.properties`

3. Edit this file with a text editor and set its content to the following five lines:
extension_uri=/jakarta/isapi_redirect.dll
log_file=[maestro_install_folder]\logs\isapi_redirect.log
log_level=info
worker_file=[connector_folder]\workers.properties
worker_mount_file=[connector_folder]\uriworkermap.properties
where you replace [maestro_install_folder] and [connector_folder]
(two occurrences of the latter) with the correct values for your system.

32.2.1 Configuring IIS 6

1. Open the IIS Manager.
2. In the Manager, open the node for the local computer.
3. Right-click on the **Web Service Extension** node.
4. Select **Add a new Web service Extension** from the popup menu. In the dialog that opens:
 - a. Enter a meaningful name for the extension (for example "LISTSERV Maestro").
 - b. Add one required file: Click the **[Add]** button, then the **[Browse]** button, then select the file [connector_folder]\isapi_redirect.dll
 - c. Check the **Set extension status to Allowed** checkbox.
 - d. Click **[OK]**.
5. Still in the Manager, open the **Web Sites** node.
6. Determine the web site that you want to use to act as a "proxy" for LISTSERV Maestro.
7. Right click on the node for this web site.
8. Select **New > Virtual Directory** from the popup menu. In the wizard, create a new virtual directory with the following settings:
Alias: Must be "jakarta" (without the quotes).
Path: Select the [connector_folder]
Permissions: Grant both "Read" and "Execute".
9. Again, in the Manager, right click on the node for this web site.
10. Select **Properties** from the popup menu. In the properties dialog:
 - a. Go to the **ISAPI Filters** tab.
 - b. Click **[Add]**.
 - c. Enter a meaningful filter name (for example "LISTSERV Maestro").
 - d. Click **[Browse]** to select the executable. Select the file [connector_folder]\isapi_redirect.dll

32.2.2 Configuring IIS 7



Note: At this time, we have not been able to test the following procedure ourselves, so it is described to the best of our knowledge.

In the Server Manager:

1. Both "ISAP Filters" and "ISAPI Extensions" must be installed as role services.
2. Determine the web site that you want to use to act as a "proxy" for LISTSERV Maestro.
3. Right click on the node for this web site.
4. Select **Add Application** from the popup menu. In the dialog, enter the following settings:
Alias: Must be "jakarta" (without the quotes).
Physical path: Select the [connector_folder].
5. Select the web site node again, then in the right pane, double click the "ISAPI Filters" symbol, then right click in the right pane.
6. Select **[Add]** from the popup menu. In the dialog:
 - a. Give the filter a meaningful name (for example "LISTSERV Maestro").
 - b. For the executable, select the file
[connector_folder]\isapi_redirect.dll
7. Select the web site node again, then in the right pane, double click the "Handler Mappings" symbol, then right click in the right pane.
8. Select **Add Module Mapping** from the popup menu. In the dialog:
 - a. **Request path:** Must be "*.dll" (without the quotes).
 - b. **Module:** IsapiModule
 - c. **Executable:** Select the file
[connector_folder]\isapi_redirect.dll
 - d. **Name:** Specify a meaningful name (for example "LISTSERV Maestro").
 - e. After you click **[OK]** in the dialog, you are asked if you want to allow this ISAPI extension. You must answer with **Yes**.

32.3 Completing the Configuration

After the configuration steps above have been completed, you must restart LISTSERV Maestro and IIS to activate the changes you made.

If everything is configured correctly, you should then be able to access LISTSERV Maestro through IIS. For example, if the web site for which you configured the proxying above is called "www.mycorp.com", then you can now access LISTSERV Maestro with the following URL:

`http://www.mycorp.com/lui`

As a last step, you should now log in to the Administration Hub and configure the access URLs to reflect that access to LISTSERV Maestro is now via the default port, i.e. the access URLs must no longer include the port element.

For example, if your LISTSERV Maestro is running on port 8080, then before you performed the above proxy configuration, your access URLs will look something like "http://SERVERNAME:8080". You now need to change these to plain URLs without a port element, like "http://SERVERNAME".

If you are configuring a fresh installation of LISTSERV Maestro, then that is all you need to do. But if you are configuring a LISTSERV Maestro instance that has already been in use for a while with the non-standard port (for example, 8080), then people might already have bookmarks to the LISTSERV Maestro access pages with this non-standard port.

After you have now changed the access URLs to no longer contain the non-standard port, users will no longer be allowed to log in to LISTSERV Maestro if they access it via these bookmarks. You should therefore set the old access URLs (the ones with the non-standard port) as access URL aliases. With such an alias, if a user accesses LISTSERV Maestro with the old URL (through an old bookmark), then they will still be allowed to login, and after login, they are automatically redirected to the new URL (with the standard port).

Glossary of Terms

A

Action Tracking – A cookie-based method to track actions and page visits that recipients perform after they have read your email message and followed the links to the target website.

Administration Hub (HUB) – A component of the LISTSERV Maestro program that allows the administrator to create user accounts, and assign and change settings for the entire application.

AOL Rich Text Formatting – A message formatting style specific to the America Online email client. Combines HTML elements with text elements. Users must have an AOL email client to properly view this format. This setting is obsolete because recent versions of the AOL email client accept HTML emails and have actually stopped accepting the AOL Rich Text format.

B

Bounce – An Email message that is returned as undeliverable

Bounce Server – An optional dedicated LISTSERV server used exclusively to process bounced email.

C

Click-Through Event – A trackable occurrence available with text and HTML email messages that records each time a URL contained in the message is clicked.

Client System – A computer used to access the Maestro HUB or LUI components.

Column – A vertical set of data, as in a table or spreadsheet.

Content Template – Used to automatically define a job's content, and allows a user to create professional-looking HTML messages without any HTML coding. An email job's content definition can be based on these ready-to-use message templates that contain placeholders that, when selected in the template designer, will let you fill in predefined areas with your own text or images. In addition, all text, including any changes you may make down the road, will be entered into both the HTML and the text part of your message so you only need to enter the text once. Each placeholder can represent one or several instances of plain text or HTML, or a linked or inline binary. When using the template designer, you can also preview each placeholder, providing a full picture of what the template will look like when it is finished.

D

Data Warehouse – Maestro's Data Warehouse stores and manages recipient profiles and target groups stored within LISTSERV Maestro.

Database – A large collection of data organized with inter-related data tables for rapid search and retrieval, managed as an entity by a DBMS.

Database Client – Software used to access a database server.

Database Driver – A program installed on a workstation or server to allow programs on that system to interact with a DBMS.

Database Plugin – A feature that allows LISTSERV Maestro to connect to a driver for a particular DBMS.

Database Server – A single server running a DBMS to manage one or more databases.

Database Server Name – Upon installation, LISTSERV Maestro will automatically create a Database Server Name for the system database. This name is viewable through the Administration Hub under Connection Settings for LISTSERV Hosted Lists, and must be entered into LISTSERV's site configuration file.

Datasets -- A recipient dataset is a collection of data organized into fields and pertaining to recipients. The fields making up a dataset can have different types of properties that determine the kind of data within them, such as text, numbers, menu selections, dates, and so on. The data administrator designs the datasets within a recipient warehouse, defining each field and the type of data it holds. The recipient dataset also contains mailing lists created by the data administrator that use the recipient data for job definition. The data is shared across all mailing lists that are created within the dataset. Individual lists within the dataset are allowed to have additional fields of data that pertain just to those lists.

DBMS – Database Management System (DBMS) is a software product for the management of databases. Examples are: DB2, MySQL, Oracle, MS SQL Server.

Delimiter – The character or symbol that is used to separate one item from another. In text files imported into databases, commas are often used as delimiters. A delimiter is the same as a separator character.

DISTRIBUTE Job – A DISTRIBUTE job is a specially-formatted message sent to LISTSERV from LISTSERV Maestro. The DISTRIBUTE job contains an email message and recipient data.

Drop-In – A drop-in allows for content to be pulled from some source and inserted into a message sent by LISTSERV Maestro. For instance, an unsubscribe banner could be automatically added to an outgoing message.

DSN – Data Source Name (DSN) is required in the setup of ODBC in order to specify the connection information for a database server. Database clients use the information contained within the DSN to locate and log on to a database.

E

Email Job – In LISTSERV Maestro an email job is the creation of a customized list of recipients matched with a customized email message that is scheduled for delivery and then sent out.

Email Merge – Placing variables that are extracted from a database into an email message template. This operation permits individual personalization of otherwise bulk email messages.

Encoding – Is the transformation of data into digital form. With text encoding, different character sets encode text files differently based on language and other variables. If a special character set was used to encode a text file, that same encoding scheme needs to be used to interpret the data correctly. LISTSERV[®] Maestro allows for the selection of encoding based upon the original encoding scheme of the uploaded text file. For example, if special encoding was used to initially create (and save) the text file (e.g. ISO-7 encoding for a file with Greek characters, or a Unicode encoding), the same encoding will have to be selected in LISTSERV[®] Maestro so that the system interprets the uploaded data in the same way it was saved.

External System Database – Maestro is installed with a built-in internal MySQL database. An alternate external database may be installed and configured to take the place of the internal one.

F

G

H

Header – A special row of data that defines and labels the columns in a text file.

Host – Refers to a computer system on which one or more of the Maestro components resides. The “Hostname” is the name of host system (e.g., MAESTRO.EXAMPLE.ORG).

Hosted Lists -- Hosted lists contain data from the dataset. They can also have their own data fields that are not shared among lists in the dataset, but belong exclusively to the list. Lists that do have their own fields will also have their own web subscription forms generated when the list is created. All the fields that appear in the dataset and in a particular list can be used as merged fields for messages sent to that list.

Hosted LISTSERV List – On some platforms, it is possible for LISTSERV Maestro to create traditional LISTSERV lists and store the subscription data for these lists such that they are accessible from the LISTSERV Maestro subscription pages. These types of

lists are referred to as Hosted LISTSERV Lists or HLLs because their data is “hosted” within LISTSERV Maestro’s system database.

Hosted Recipient Data – A collection of data organized in columns and rows related recipients and stored inside LISTSERV Maestro.

Hosted Recipient List – Lists that are controlled completely by LISTSERV Maestro.

HTTP – Hyper Text Transfer Protocol (HTTP) is the language used by web browsers and web servers to communicate with each other.

HTTPS – Secure HTTP. Similar to HTTP, but the communication is encrypted, making it more difficult for a third party to eavesdrop on the communication.

HUB – See “Administration Hub”.

I

Identity – A collection of several accounts that belong to one and the same "identity", usually a person. By collecting all accounts of one person into an identity, LISTSERV Maestro knows that these accounts all belong together. As a result, the user is then allowed to switch between the accounts in the identity without having to perform an actual logout and login. In other words, if a user logs in with one account that belongs to an identity, then this user can switch over to all other accounts in the same identity without having to first logout the old account and then login again with the new account.

Internal System Database – Maestro is installed with a built-in MySQL database. This internal database is automatically configured upon installation. An alternate external database may be installed and configured instead.

J

Java – The Java programming language and runtime environment are the technology on which LISTSERV Maestro is built.

JDBC – Java Database Connector (JDBC) allows Java applications such as LISTSERV Maestro to connect to database servers.

Job ID – A unique identifier assigned to each mailing sent through LISTSERV Maestro. May have a Job ID Prefix attached to it.

K

Keystore – In order for SSL to work, the client must have a set of “keys” that contain the digital signatures of trusted servers. These keystores allow for the client to verify the identity of a trusted server.

L

List Archive – A LISTSERV-managed archive containing all messages sent to a Hosted LISTSERV List.

LISTSERV® – An application that allows users to create and maintain email lists on their corporate networks or on the Internet. LISTSERV supports all types of email lists: newsletters (moderated and unmoderated), discussion groups, and direct marketing campaigns. List sizes can range from a few participants in a discussion group to several million in a newsletter. Every list and its archives can be maintained through a simple web interface, which can be fully customized to match a website profile. When used within LISTSERV Maestro, LISTSERV receives email jobs from LISTSERV Maestro and prepares them for delivery. It is also used to process bounced mail for LISTSERV Maestro mailings. Additionally, LISTSERV may act as an interface between LISTSERV Maestro and an external DBMS. When Hosted LISTSERV Lists (HLLs) are used, LISTSERV Maestro acts as the “DBMS back-end” to traditional LISTSERV lists and also provides an interface for management of subscriber data for the LISTSERV lists.

LISTSERV Maestro – The software suite comprised of the Administration Hub, Maestro User Interface, and Maestro Tracker.

Lookup Table – A set of values that is used for the values in a selection menu. Lookup tables are shared across a recipient warehouse so multiple datasets can use them

LUI – See “Maestro User Interface”.

M

Maestro Tracker (TRK) – A component of the LISTSERV Maestro program that receives and compiles tracking data from delivered email messages.

Maestro User Interface (LUI) – A component of the LISTSERV Maestro program that allows regular users to create email jobs and tracking reports.

N**O**

ODBC – Open Database Connectivity (ODBC) is the means by which LISTSERV on Windows connects to databases. LISTSERV Maestro may also use an ODBC plugin for its database connectivity for read-only access to external recipient data.

Open-Up Event – A trackable occurrence available with HTML email messages that records each time a message is opened by a recipient. Tracking is dependent on the willingness of the recipients to be tracked; therefore, open-up counts are usually lower than the actual number of open-up events.

P

Port Number – A port number is a number assigned to a particular network service on a host. For example, SMTP usually uses port 25, while HTTP is usually port 80.

POSTMASTER – Used generally, a “postmaster” is someone responsible for the administration of an Email server. In LISTSERV, the POSTMASTER site configuration parameter specifies the email addresses of individuals who have administrative control over LISTSERV (and thus may create or delete lists, send DISTRIBUTE jobs, etc.).

Q

Quote character – In a SQL statement: a character (usually the single quote) used to enclose string literals, to set them off from the rest of the SQL statement.

In a text file (CSV-file) containing data: a character or symbol used to surround the value of a column if the value contains the separator character in the actual data. This is necessary to ensure that the appearance of the separator character in the data is not interpreted as an actual separation. For example, if a comma (,) is used as the separator character in a database file, all the fields of data are separated by a comma. If the comma is also used within a field, the quote character must be used to surround the entire field. If the quote character is used within a field, it must be doubled, or “escaped.”

R

Recipient Profile – Data stored within LISTSERV Maestro particular to a unique recipient. Contains at least an Email address, but may also contain other user data such as name, mailing preferences, etc.

Recipient Warehouse – The repository for a group's data including lookup tables, datasets, hosted lists, and recipient data.

RFC – Request for Comments (RFC) are documents that explain the rules that email and other software products must follow in order to work cooperatively with each other on the Internet. Understanding the rules is often helpful for understanding and troubleshooting problems.

S

Select Statement – A SQL statement in form of a query that is issued to a database to retrieve data.

Separator Character – A character or symbol used to separate one item from another. In text files exported from databases, commas are often used as separator characters. A separator character is the same as a delimiter.

SMTP – Simple Mail Transfer Protocol (SMTP) is the protocol used by servers that send and receive email messages over the Internet. An SMTP server is a mail transfer agent (MTA) that uses the SMTP protocol to talk to other MTAs.

SMTP Worker Pool – Set on the LISTSERV Connection screen, this feature lets you specify a LISTSERV worker pool to use for specific delivery situations. You can specify different worker pools for standard deliveries and for test deliveries.

SQL – Structured Query Language (SQL) is a standardized query language for requesting information from a database.

SQL Statement – A statement written in SQL that is issued to a database to retrieve data or to create, insert, update, or delete data in the database.

SSL – Secure Socket Layer (SSL) is the means by which secure communications (such as HTTPS transactions) are encrypted.

System Database – The database in which the LISTSERV Maestro User Interface (LUI) stores its system data (e.g., tracking data, job data, hosted recipient data, etc.) When set up as an external system database, it may or may not be managed by the same DBMS server as additional User Databases.

T

Tablespace – The digital “space” within a database allotted to a particular user or set of tables.

Target Groups – Predefined recipient lists, complete with name and description, created by the data administrator. Target groups can simplify and streamline the use of data sources, including databases, uploaded text files, and email lists, to select recipients and recipient data to the point where end users do not need to know anything about how and where data is stored.

TCPGUI – The protocol used by Maestro to communicate with LISTSERV. The default port for TCPGUI is 2036.

Tracking URL – Defines the URL for the Maestro Tracker Server.

TRK – See “Maestro Tracker”.

Trusted Root Certificate – Certain “root” agencies serve as registration storehouses for digital “keys”. The root certificates (shipped with most web browsers) contain the digital signatures of the root agencies.

U

User Database – An external database from which LISTSERV Maestro retrieves recipient data. For example, LISTSERV Maestro can be used to send mailings to email addresses extracted directly from an organization’s customer database, unlike other products that *require* you to store the recipient data in their own database. May or may not be on the same DBMS server as the External System Database.

V

W

X

Y

Z

Index

A

- Accessing
 - disallowing with same user account 165
 - HUB 6
 - log files 145
 - restricting access to components 163
 - restricting login attempts 168
- Action Tracking
 - enabling
 - default 42
 - group 73
 - user 73
- Admin
 - emergency HUB access 269
- Admin Password
 - changing 15
- administer target groups 65
- Administration
 - configuring backups 51
 - email notifications 54
 - introduction 1
 - policies 51
 - refreshing the Subscriber Page
 - translations 57
 - runtime administration 51
 - special administrative user account 81
 - system shutdown 51
- Administrative Policies 51
 - configuring backups 51
 - email notifications 54
 - user restrictions 53
- Archived Jobs
 - importing 102
- Archiving
 - completed jobs 101
 - delivered jobs 101
- Auto-Archive Settings
 - defining for an account or group 76
 - defining the default settings 46

B

- Backups
 - configuring 51
 - configuring a backup time 131
 - configuring external post-backup processes 132
 - configuring the backup history 134
 - configuring the backup location 133

- creating a Test-Bed Backup on the original system 143
 - ID 135
 - restoring 131, 135
 - restoring a Test-Bed Backup into the Test System 144
 - saving 131
 - saving to an external medium 134
 - using a Test-Bed Backup 141
- Build Version
 - finding 5

C

- Character Sets
 - allowing bi-directional charsets 254
 - defining the default mail charset 253
 - disallowing bi-directional charsets 254
 - using
- Configuring
 - backups 51
 - LISTSERV Maestro for first use 13
 - port usage 150
 - the application server shutdown port 155
 - the backup history 134
 - the backup location 133
 - the backup time 131
 - the external database 124
 - the external post-backup processes 132
 - the HTTP port 150
 - the HTTPS port 153
 - the internal communication port 153
 - the internal database connection port 154
 - the tracker communications port 154
- Content Restrictions
 - defining for an account or group 70
 - defining the default settings 39
- Content Templates
 - changing ownership 107
- Current Build Version
 - finding 5

D

- Dashboard 90
 - adding reports 97
 - copying 98
 - creating from scratch 99
 - creating multiple 98

- Current and Upcoming Deliveries
 - section 91
 - Currently in the System section 90
 - deleting 100
 - hiding/showing sections 93
 - Jobs Due Next section 91
 - managing reports 96
 - Recent Deliveries section 91
 - refreshing the data 92
 - renaming 92
 - renaming reports 97
 - re-organizing the reports 97
 - switching 99
 - System Metrics section 90
 - viewing additional report details 97
 - viewing system metrics 93
- Database Plugins 116
- IBM DB2 v7.2 117
 - IBM DB2 v8.2 116
 - MySQL ConnectorJ Driver 117
 - MySQL L-Soft Driver 118
 - ODBC Driver 120
 - Oracle 8i, 9i, 10g Thin Driver 118
 - registering 122
 - SQL Server i-net SPRINTA Driver 120
 - SQL Server jTDS Driver 119
 - SQL Server Microsoft Driver 119
- Dataset
- deleting a Hosted List 114
- Datasets
- changing ownership 112
 - deleting 112, 114
- Default URL Settings
- defining 34
 - default HUB access URL for users 35
 - default LUI access URL aliases for users 35
 - default LUI access URL for admin 36
 - default LUI access URL for users 35
 - default tracking URL 35
 - HUB access URL for admin 36
- Disk Space
- monitoring 93
- Distributed Components
- configuring LISTSERV Maestro
 - components with server name aliases or proxies 234
 - fresh installation 226
 - introduction 225
 - moving components to another server 226
 - moving the Administration Hub component 229
 - moving the database external component 232
 - moving the Maestro User Interface component 226
 - moving the Tracker component 230
 - server name aliases and proxies 232
- DomainKey Settings
- defining for an account or group 75
- DomainKeys
- using 201
- DomainKeys Settings
- defining the default settings 45
- Drop-Ins
- changing ownership 107
- ## E
- Email Jobs
- limiting the number sent
 - as the system default 40
 - for accounts 71
- Email Notifications
- administration 54
 - defining different SMTP servers 54
 - sending an email notification after each backup 54
 - testing 56
- Evaluation Mode 237
- External Access Settings
- defining for an account or group 76
- External Database
- adding
 - configuring 124
 - defining connections 115
 - removing
- External Database See *System Database*
- ## F
- Firewall
- install behind 159
- ## G
- Global Settings
- User Interface 29
 - Database Plugins 31
 - default 33
 - Default Auto-Archive Settings 46
 - Default Content Restrictions 39
 - Default DomainKeys Settings 45
 - Default Hosted Data Import Restrictions 43
 - Default Hosted Data Settings 42
 - Default LISTSERV Connection 36
 - Default Recipients Restrictions 40
 - Default Size Limits 39
 - Default Tracking Restrictions 42
 - Default URL Settings 34

- General Administration 29
 - LISTSERV Web Interface Access 33
 - System Database Connection 32
 - User Rights 32
- Groups
 - assigning 60
 - assigning user rights 62
 - deleting 77
 - settings 64
 - auto-archive settings 76
 - content restrictions 70
 - DomainKey settings 75
 - external access settings 76
 - hosted data settings 73
 - job ID prefix 69
 - LISTSERV connection 66
 - recipient restrictions 71
 - size limits 69
 - tracking restrictions 73
 - URL settings 67
 - user rights 64
- H**
- Hosted Data Settings
 - defining for an account or group 73
 - defining the default settings 42
- Hosted Lists
 - deleting from a dataset 114
- HTML Upload Applet
 - Microsoft Java Virtual Machine 260
 - Sun Java Plugin 259
 - updating 259
- HTTP Port
 - configuring 150
- HTTP Ports
 - configuring 150
 - defining 157
- HTTPS Port
 - configuring 153
- HUB Access
 - emergency Admin access 269
- HUB Interface
 - accessing 6
 - changing a password 15
 - editing the INI file 195
 - introduction 8
 - login 6
 - ports 149
 - Single Sign-on 6
 - switching between Interface 7
- I**
- Icons
 - Help 11, 89
- Identities 59
 - creating 61
 - deleting 77
 - editing 60
- IIS
 - sharing a server with 271
- Importing
 - archived jobs 102
- INI Files
 - Administration HUB entries 195
 - editing 189
 - Maestro User Interface entries 190
 - Tomcat entries 197
 - Tracker entries 196
- Installation
 - moving LISTSERV Maestro from one server to another 217
- Integration
 - between LISTSERV and LISTSERV Maestro 205
- Interface Links
 - between Membership Area and Subscriber's Corner 213
 - defining 205
- Internal Database
 - adding
 - removing
- Internal Database See *System Database*
- International Character Sets
 - See *Character Sets*
- IP Addresses
 - defining 157
- J**
- Job ID Prefix
 - setting 69
- Jobs
 - archiving 101
 - changing ownership 103
 - importing archived jobs 102
- K**
- L**
- Links
 - defining between LISTSERV and LISTSERV Maestro 205
 - defining between Membership Area and Subscriber's Corner 213
- ListPlex
 - on the Tomcat Server 246

- LISTSERV
 - preparing for LISTSERV Maestro 17
 - preparing Maestro to send distribute jobs 25
 - preparing to allow HLL 19
 - preparing to process distribute jobs from Maestro 17
 - specifying a separate LISTSERV instance for processing bounces 27
 - specifying the LISTSERV host with different internal and external names 26
 - using existing lists with Maestro 28 with LISTSERV Maestro 17
- LISTSERV Connection
 - defining for an account or group 66
 - defining the default settings 36
- LISTSERV Maestro
 - configuring for first use 13
 - configuring to bind to different ports on different IP addresses 157
 - default settings 33
 - defining links between Membership Area and Subscriber's Corner 213
 - defining links with LISTSERV 205
 - editing INI files 189
 - enabling single sign-on 209
 - in evaluation mode 237
 - installing behind a firewall 159
 - integration with LISTSERV 205
 - introduction 1
 - LISTSERV instances 17
 - log files 145
 - making aware of the server certificate 178
 - mixing SSL and non-SSL access on one server 181
 - moving an installation 217
 - preparing LISTSERV for 17
 - preparing LISTSERV to allow HLL 19
 - preparing LISTSERV to process distribute jobs 17
 - preparing to send distribute jobs to LISTSERV 25
 - securing a server with multiple host names 180
 - securing access with SSL 169
 - setting the default tracking URL 46
 - settings for the User Interface 29
 - sharing a server with IIS 271
 - specifying a separate LISTSERV instance for processing bounces 27
 - standard default ports 257
 - User Roles 2
 - using existing LISTSERV lists 28
- LISTSERV Maestro Lite 263
 - admin restrictions 267
 - restrictions 264
 - user restrictions 264
- LISTSERV nodes
 - serving multiple nodes on a single server 246
- LISTSERV Web Interface
 - adding to the Tomcat server 239
- Log Files
 - accessing 145
 - Remote Log Access 145
 - Subscriber Activity Change Log 146
- Login
 - auto login 48
 - HUB 6
- Login Restrictions
 - setting for an account or group 67
- Logs
 - remote access 3
- Lookup Tables
 - changing ownership 112
 - deleting 112, 113
- M**
- Maestro Tracker
 - configuring the Tracker
 - communications port 154
 - editing the INI file 196
 - ports 149
- Mixed Mode
 - mixing SSL and non-SSL access on one server 181
- Moving
 - LISTSERV Maestro installation from one server to another 217
- N**
- O**
- P**
- Password
 - changing 15
- Ports
 - configuring the application server shutdown port 155
 - configuring the HTTP Port 150
 - configuring the HTTPS port 153
 - configuring the internal communication port 153
 - configuring the internal database connection port 154

- configuring the tracker
 - communications port [154](#)
 - configuring usage [150](#)
 - for the HUB [149](#)
 - for the TRK [149](#)
 - for the User Interface [149](#)
 - standard defaults [257](#)
 - using non-standard ports [149](#)
- Q**
- R**
- Recipient Datasets
 - changing ownership [112](#)
 - deleting [112](#)
- Recipient Profiles [185](#)
- Recipient Restrictions
 - defining [71](#)
 - defining the default settings [40](#)
- Recipient Warehouse
 - changing ownership [108](#)
- Remote Log Access [3](#)
- Remote Version Query
 - HUB [5](#)
 - LUI [5](#)
 - TRK [5](#)
- Reports
 - adding to the Dashboard [97](#)
 - changing ownership [103](#)
 - managing on the Dashboard [96](#)
 - monitoring daily usage [94](#)
 - monitoring status [94](#)
 - renaming on the Dashboard [97](#)
 - re-organizing on the Dashboard [97](#)
 - viewing details on the Dashboard [97](#)
- Restrictions
 - disallowing concurrent access with
 - the same user account [165](#)
 - for login attempts [168](#)
 - IP Address [163](#)
 - to components [163](#)
- Runtime Administration
 - settings [51](#)
- S**
- Secure Communication
 - components [172](#)
 - creating an unsigned server certificate [174](#)
 - installing a server certificate [173](#)
 - installing a signed server certification [176](#)
 - installing a trusted root certificate [177](#)
 - introduction [169](#)
 - making Maestro aware of the server
 - certificate [178](#)
 - mixing SSL and non-SSL access on
 - one server [181](#)
 - performing a CSR [175](#)
 - securing a server with multiple
 - host names [180](#)
 - securing access with SSL [169](#)
 - securing the Trusted Root Certificate
 - Keystore [173](#)
- Sender Profiles
 - changing ownership [107](#)
- Server
 - sharing with IIS [271](#)
- Server Name Aliases and Proxies [232](#)
- Single Sign-On
 - enabling [209](#)
- Single Sign-on [6](#)
 - switching to LUI [7](#)
- Size Limits
 - defining the default [39](#)
 - setting for an account or group [69](#)
- SLL Cipher
 - supported [183](#)
- SMTP Workers
 - defining separate for deliveries [37](#)
- SSL
 - securing access [169](#)
- Standard Delivery
 - defining separate SMTP workers [37](#)
- Subscriber Pages
 - refreshing the translations [57](#)
- System Database
 - adding an external DB
 - available plugins [116](#)
 - configuring the external DB [124](#)
 - DB2 [128](#)
 - defining external database
 - connections [115](#)
 - introduction [123](#)
 - MySQL [129](#)
 - optimization [130](#)
 - Oracle [127](#)
 - preparing [126](#)
 - removing an internal DB
 - SQL Server [127](#)
- System Metrics
 - defining daily system update time [30](#)
 - defining send email after each update [30](#)
 - monitoring disk space [93](#)
 - on the Dashboard [93](#)
 - reports [94](#)
- System Requirements
 - Client [5](#)
- System Shutdown [51](#)

T

- Target Group Categories
 - changing ownership 110
 - deleting 112
- Target Groups
 - changing ownership 110
 - deleting 112
- Test Deliveries
 - defining separate SMTP workers 37
- Test-Bed Backups 141
 - creating 143
 - restoring into the Test System 144
- Tomcat
 - adding content 239
 - adding content as a new context 240
 - adding the LISTSERV Web Interface 239
 - defining the default context 242
 - editing the INI file 197
 - enabling access logging for added content 244
 - enabling access logging for custom content 245
 - enabling access logging for WA 244
- Tomcat Server
 - ListPlex 246
 - serving multiple LISTSERV nodes on a single server 246
- Toolbar 8, 81
 - Accounts and Identities menu 8
 - Back To menu 88
 - behavioral changes in the User Interface 90
 - Dashboard menu 85
 - Dataset menu 85
 - Global Settings menu 10
 - Group menu 9
 - Help icon 11, 89
 - Hosted List menu 86
 - Linked LISTSERV List menu 87
 - LISTSERV menu 88
 - Logout menu 89
 - Logout option 11
 - Lookup Table menu 87
 - Preferences menu 89
 - Recipient Warehouse menu 83
 - Report menu 83
 - User Account menu 8
 - Utility menu 11, 84
- Tracker
 - configuring the communications port 154
 - editing the INI file 196
- Tracking 185

Tracking Restrictions

- defining for an account or group 73
- defining the default settings 42

U

URL Settings

- defining for an account or group 67
- defining the default settings 34
- HUB Access URL 69
- LUI Access URL 67
- LUI Access URL Aliases 68
- Subscriber Pages Access URL 68
- Subscriber Pages Access URL Aliases 68
- Tracking URL 69

User Accounts 59

- assigning group user settings 62
- assigning single user settings 62, 63
- creating 60
- deleting 77
- disallowing concurrent access 165
- editing 60, 63, 64
- settings 64

- auto-archive settings 76
- content restrictions 70
- DomainKey settings 75
- external access settings 76
- hosted data settings 73
- job ID prefix 69
- LISTSERV connection 66
- recipient restrictions 71
- size limits 69
- tracking restrictions 73
- URL settings 67
- user rights 64
- special administrative account 81

User Interface

- editing the INI file 190
- ports 149
- setting the default tracking URL 46
- settings 29

User Restrictions 53

User Right Settings

- administer content templates 65
- administer hosted recipient datasets 65
- create new drop-in content elements 65
- create new jobs 64
- create new reports 65
- create new sender profiles 65
- defining for an account or group 64
- defining the default settings 32
- link recipient datasets to LISTSERV 65

- User Rights
 - assigning [62](#)
 - managing [62](#)
- User Roles
 - Data Warehouse Administrator [2](#)
 - Database Administrator [2](#)
 - LISTSERV Maestro Administrator [2](#)
 - LISTSERV Site Administrator [2](#)
 - Maestro User [2](#)
 - SMTP Server Administrator [2](#)
 - System Administrator [2](#)
- User Settings
 - default settings [33](#)
- Using Rights
 - editing [62](#)

V

W

- What's New
 - configuring aliases for access URLs [xvii](#)
 - configuring LISTSERV Maestro to bind to different HTTP Ports on Different IP addresses [xvii](#)
 - defining the hosted data import restrictions [xv](#)
 - emergency admin HUB access [xvii](#)
 - enabling action tracking [xv](#)
 - enabling job data export [xv](#)
 - hosted data changes [xvi](#)
 - mixing SSL and non-SSL access on one server (mixed mode) [xvi](#)
 - moving a LISTSERV Maestro installation [xvii](#)
 - new look and feel [xv](#)
 - sharing a server with IIS [xvii](#)
 - SSL cipher support [xvi](#)
 - switching between interfaces [xv](#)
 - system metrics
 - monitoring [xv](#)
 - Toolbar changes in the User Interface [xvi](#)
 - URL settings [xvi](#)

X

Y

Z

