

L-Soft Sweden AB



Administrator's Manual

LISTSERV[®] Maestro, version 2.1



Last Updated: 6/27/2006 2:35 PM

This page left intentionally blank

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 International, Inc. 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, provided that the copies are not used for commercial advantage, that the source is cited and that the present copyright notice is included in all copies, so that the 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 © 2006 L-Soft Sweden AB
All Rights Reserved Worldwide.

L-SOFT, LISTSERV, ListPlex, and LSMTP are registered trademarks of L-Soft international, Inc. LMail is a trademark of L-Soft international, Inc.

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.

MySQL is a registered trademark of MySQL AB in the United States, the European Union, and other countries.

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 for LISTSERV are available in ASCII-text format from LISTSERV and in PDF format from **ftp.lsoft.com**. They are also available on the World Wide Web at the following URL:

URL: <http://www.lsoft.com/manuals/index.html>

L-Soft invites comment on its manuals. Please feel free to send your comments by email to: MANUALS@LSOFT.COM

Table of Contents

Preface - About This Manual	ix
What's New in LISTSERV® Maestro 2.1	x
Subscriber Interface Enhancements	x
The Subscriber Interface Customization Wizard	x
New Subscriber Options	xi
New Resource Translation Tool	xi
Recipient Data Management Enhancements	xi
Automatic Bounce Processing	xii
Flexible Dataset Administration	xii
Shared Subscriber Management and Reporting	xii
Tracking Enhancements	xiii
Tracking Restrictions	xiii
Tracking Links in Drop-Ins	xiii
Administration Enhancements	xiii
LUI Admin Account	xiii
Refreshing the Subscriber Page Translations	xiii
Creating a Test-Bed Backup	xiv
DomainKeys Signatures	xiv
System Support Enhancements	xiv
Section 1 Introduction to LISTSERV® Maestro Administration	1
1.1 Maestro User Roles	2
1.2 Remote Version Query	3
1.3 Client System Requirements	3
1.4 Accessing the Hub Administrator Interface	4
1.5 Understanding the Hub Administrator Interface	4
Section 2 Configuring LISTSERV Maestro for First Use	9
Section 3 Changing the Administrator Password	11
Section 4 The System Database	12
4.1 Configuring the External System Database	13
4.2 Preparing the System Database	15
4.2.1 General System Database Preparation	15
4.2.2 Preparing SQL Server as the System Database	16
4.2.3 Preparing Oracle as the System Database	16
4.2.4 Preparing DB2 as the System Database	17
4.2.5 Preparing MySQL as the System Database	17
4.3 General Optimization Hints for the System Database	18
4.4 Removing and Adding the Internal Database	19

Section 5 Defining External Database Connections	20
5.1 Available Database Plugins	21
5.1.1 The IBM DB2 V8.2 Thin Driver Database Plugin	21
5.1.2 The IBM DB2 V7.2 Native Driver Database Plugin	21
5.1.3 The MySQL ConnectorJ Driver Database Plugin	22
5.1.4 The MySQL L-Soft Driver Database Plugin	22
5.1.5 The Oracle 8i, 9i, and 10g Thin Driver Database Plugin	23
5.1.6 The SQL Server jTDS Driver Database Plugin	23
5.1.7 The SQL Server Microsoft Driver Database Plugin	24
5.1.8 The SQL Server i-net SPRINTA Driver Database Plugin	24
5.1.9 The ODBC Driver Database Plugin	25
5.2 Registering a Database Plugin	26
Section 6 LISTSERV and LISTSERV Maestro	28
6.1 Preparing LISTSERV for LISTSERV Maestro	28
6.1.1 Preparing LISTSERV to Process DISTRIBUTE Jobs from LISTSERV Maestro	28
6.1.2 Preparing LISTSERV to Allow Maestro Hosted LISTSERV Lists	30
6.1.3 Preparing LISTSERV for Database Access	32
6.2 Preparing LISTSERV Maestro to Send DISTRIBUTE Jobs to LISTSERV	36
6.2.1 Specifying the LISTSERV Host with Different Internal and External Names	37
6.2.2 Specifying a Separate LISTSERV Instance for Processing Bounces	38
6.3 Using Existing Lists with LISTSERV Maestro	38
Section 7 Settings for the Maestro User Interface	40
7.1 Application Settings	40
7.2 Application Default Settings	43
7.3 Setting the Default Tracking URL	49
7.3.1 Multiple Tracking URLs	50
Section 8 Administrative Policies	51
8.1 Configuring Backups	51
8.2 Runtime Administration and System Shutdown	51
8.3 User Restrictions	52
8.4 Administrative Email Notifications	54
8.4.1 Testing Email Notifications	55
8.5 Refreshing the Subscriber Page Translations	56
Section 9 Creating and Administering User Accounts	57
9.1 Creating a New User Account	57
9.2 Editing Account Information and Assigning Single User Settings	58
9.2.2 Editing General User Settings	58
9.2.3 Editing Component Specific Settings for Single and Group Users	59
9.3 Deleting a User Account	67

Section 10 Special Administrative User Account	68
10.1 Archiving Delivered Jobs	68
10.2 Importing Archived Jobs	70
10.3 Changing Job and Report Ownership	70
10.4 Changing Sender Profile and Drop-In Content Element Ownership	74
10.5 Changing Recipient Dataset and Target Group Ownership	76
10.5.1 Changing Ownership of Recipient Datasets and Lookup Tables.....	76
10.5.2 Changing Ownership of Recipient Target Groups	79
Section 11 Saving and Restoring a Backup	80
11.1 Configuring the Backup Time.....	80
11.2 Configuring External Post-Backup Processes	81
11.3 Configuring the Backup Location	82
11.4 Configuring the Backup History	83
11.5 Saving a Backup to an External Medium	84
11.6 Identifying the Backup: The Backup ID	84
11.7 Restoring a Backup	84
Section 12 Using a LISTSERV Maestro Test-Bed Backup	88
12.1 Creating a Test-Bed Backup on the Original System.....	90
12.2 Restoring a Test-Bed Backup into the Test System	90
Section 13 Maestro Logs	92
13.1 Remote Log Access.....	92
13.2 Subscriber Activity Change Log.....	93
Section 14 Using Non-Standard Ports	96
14.1 Ports Used by LISTSERV Maestro	96
14.1.1 Ports used by the Administration Hub	96
14.1.2 Ports used by the Maestro User Interface	96
14.1.3 Ports used by Maestro Tracker	96
14.2 Configuring Port Usage.....	97
14.2.1 Configuring the HTTP Port	97
14.2.2 Configuring the Internal Communication Port.....	99
14.2.3 Configuring the Tracker Communications Port.....	100
14.2.4 Configuring the Internal Database Connection Port	100
14.2.5 Configuring the Application Server Shutdown Port.....	100
Section 15 Defining IP Addresses	101
Section 16 Installing Behind a Firewall	102
Section 17 Restricting Access to Components	105
17.1 IP Address Restrictions.....	105
17.2 Disallowing Concurrent Access with the Same User Account	107
Section 18 Securing Access with SSL	111

18.1 Introduction to Secure Communication	111
18.2 Which Components Should Be Secured?	114
18.3 Obtaining and Installing a Server Certificate	115
18.3.1 Securing the Trusted Root Certificate Keystore	115
18.3.2 Creating an Unsigned Server Certificate	115
18.3.3 Performing a Certificate Signing Request (CSR).....	117
18.3.4 Installing the Signed Server Certificate.....	118
18.3.5 Installing a Trusted Root Certificate.....	118
18.3.6 Making LISTSERV Maestro Aware of the Server Certificate	120
Section 19 Tracking and Recipient Profiles.....	122
Section 20 Editing LISTSERV Maestro INI Files	125
20.1 Maestro User Interface INI-File Entries.....	126
20.2 Administration Hub INI-File Entries	130
20.3 Maestro Tracker INI-File Entries	131
20.4 Tomcat INI-File Entries	132
20.4.1 Basic Tomcat Configuration Parameters	132
20.4.2 Advanced Tomcat Configuration Parameters.....	133
Section 21 Distributed Components	136
21.1 Fresh Installation with Distributed Components.....	137
21.2 Moving Components to another Server	137
21.2.1 Moving the Maestro User Interface Component to Another Server	137
21.2.2 Moving the Administration Hub Component to Another Server	139
21.2.3 Moving the Maestro Tracker Component to Another Server	141
21.2.4 Moving the Database External Component to Another Server	143
21.3 Server Name Aliases and Proxies	143
21.3.1 Configuring LISTSERV Maestro Components with Server Name Aliases or Proxies.....	145
Section 22 LISTSERV Maestro in Evaluation Mode	148
Section 23 Using International Character Sets.....	149
23.1 Defining the Default Mail Charset	149
23.2 Allowing or Disallowing Bi-Directional Character Sets	151
Section 24 Authenticating Message Origin with DomainKeys Signatures	152
Section 25 Adding Content to the Tomcat Server.....	154
25.1 Adding Content as a New Context.....	154
25.2 Defining the Default Context.....	155
Section 26 Adding the LISTSERV Web Interface to the Tomcat Server	157
Glossary of Terms.....	158
Appendix A: LISTSERV Maestro Standard Default Ports	163
Index.....	164

Figures

Figure 1 Administration Hub Home Page.....	5
Figure 2 Global Component Settings	5
Figure 3 Flow Chart of the Administration Hub.....	8
Figure 4 Change Administrator Password Screen	11
Figure 5 System Database Connection Screen	13
Figure 6 System Database Connection Details Screen for DB2	14
Figure 7 JDBC Driver Layers	25
Figure 8 ODBC Plugin Layers.....	26
Figure 9 Database Plugins	27
Figure 10 Add New Database Plugin	27
Figure 11 The Database Server Name	33
Figure 12 Using A Different External Host Name For LISTSERV	37
Figure 13 Global Component Settings for Maestro User Interface.....	40
Figure 14 General Administration of the Maestro User Interface	41
Figure 15 Database Plugins.....	42
Figure 16 System Database Connection.....	42
Figure 17 Default LISTSERV Connection	44
Figure 18 Dedicated Bounce Server	45
Figure 19 Default Tracking URL.....	45
Figure 20 Default Size Limits	46
Figure 21 Default Content Restrictions	46
Figure 22 Default Recipients Restrictions	47
Figure 23 Default Tracking Restrictions	48
Figure 24 Default Hosted Data Settings.....	48
Figure 25 Email Notifications	54
Figure 26 Refreshing the Subscriber Page Translations.....	56
Figure 27 Administer User Accounts.....	57
Figure 28 Defining User Account	58
Figure 29 Editing Account Information	58
Figure 30 User Settings for a Single (Non-Group)Member	59
Figure 31 User Settings for a Group Member	59
Figure 32 User Right Settings	60
Figure 33 Default LISTSERV Connection for a Group with Dedicated Bounce Server	61
Figure 34 LISTSERV Connection for a Single User without a Dedicated Bounce Server	62
Figure 35 LISTSERV Connection for a Group User	62
Figure 36 Tracking URL	63
Figure 37 Size Limits.....	63

Figure 38 Job ID Prefix	64
Figure 39 Drop-In Content Restrictions	64
Figure 40 Recipients Restrictions	65
Figure 41 Tracking Restrictions	66
Figure 42 Hosted Data Settings	67
Figure 43 Special Administrative User Account	68
Figure 44 Archiving a Job from the Delivered Jobs Screen	69
Figure 45 Importing a Job from the Archived Jobs Screen	70
Figure 46 Change Job Owner from Job Owner Tab	71
Figure 47 Change Report Owner Screen	72
Figure 48 Change Ownership	74
Figure 49 Select Account/Group and Item to Change Ownership	75
Figure 50 Change Drop-In Owner	75
Figure 51 Change Ownership of a Recipient Dataset or Target Group	76
Figure 52 The Recipient Datasets and Lookup Tables Screen	77
Figure 53 The Recipient Dataset and Lookup Table Administration Screen	77
Figure 54 The Recipient Dataset Administration Screen	78
Figure 55 The Hosted List Administration Screen	78
Figure 56 The Lookup Table Administration Screen	78
Figure 57 The Change Target Group Owner Screen	79
Figure 58 Changing Ownership of a Target Group Category	79
Figure 59 The General Component Settings for Administration Hub Screen	81
Figure 60 Component Communication Pathways	103
Figure 61 Multiple Logins	107
Figure 62 XML Entry for the Session Timeout	110
Figure 63 Example of Base64 Encoded Outfile	117
Figure 64 Imported Certificate	119
Figure 65 Example of Recipients Profile Data Table	123
Figure 66 Example of Recipients ID in Data Table	123
Figure 67 Sample Proxy Setup	144
Figure 68 DomainKeys Settings Screen	152

Tables

Table 1	Navigational Icons	4
Table 2	Restrictions for Hosted LISTSERV List Support	30
Table 3	Backup History	83
Table 4	Maestro User Interface INI-File Entries	126
Table 5	Administration Hub INI-File Entries	131
Table 6	Maestro Tracker INI-File Entries	131

Table 7	Basic Configuration for Tomcat INI-File Entries	132
Table 8	Advanced Configuration for Tomcat INI-File Entries	134
Table 9	Supported Charsets	149

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 by email to:

MANUALS@LSOFT.COM

The following documentation conventions have been used in this manual:

- Options, icons, and text boxes on the screen will be bold (i.e. the **Define Recipients Now** option).
- Clickable buttons will be bold and within brackets (i.e. the **[OK]** button).
- Clickable links will be bold and underlined (i.e. 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 indicated, directory paths are for Microsoft® Windows™ installations of LISTSERV Maestro. For Linux/Solaris, substitute the Maestro top-level directory for `\Program Files\L-Soft\Application Server`.



This symbol denotes an important note or warning.



This symbol denotes optional advice to save time.

What's New in LISTSERV® Maestro 2.1

LISTSERV Maestro is a highly intuitive email marketing solution that provides organizations with an advanced tool to create, send, and track opt-in personalized email messages.

LISTSERV Maestro 2.1 builds on the powerful new abilities introduced in 2.0 to collect, manage, save, and reuse data from mailing list subscribers. The Data Warehouse has become even more flexible with the addition of an interface for customizing each page in the Subscriber Web Interface. The ability to easily modify populated datasets and dataset team collaboration has made maintaining datasets less difficult and more user-friendly.

LISTSERV Maestro 2.1 continues its progress towards full integration with LISTSERV by introducing a few new features for Hosted LISTSERV Lists (HLL) and Hosted Recipient Lists (HRL). For those using Hosted LISTSERV Lists, the ability to manager subscriber options (i.e. nomail and digest) from the LISTSERV Maestro subscriber pages has been added. For those using Hosted Recipient Lists, features that were once only available for HLLs are now available for HRLs. These features are: subscriber activity change logs, automatic bounce processing, and the ability for a subscriber to set their subscription to an HRL to “nomail”.

LISTSERV Maestro 2.1 has also made improvements to the interface for the LUI administrative account, making it easier to maintain sites with many users.

These are just some of the enhancements made in LISTSERV Maestro 2.1. For a complete explanation of all changes, see the following sections.



Note: The following sections assume some familiarity with the features available in LISTSERV Maestro 2.0. It is intended as an overview of the new features for LISTSERV Maestro 2.1. For a more in-depth description about the new features in LISTSERV Maestro 2.1, refer to the User's Manual, the Administrator's Manual, and the Data Administrator's Manual.

Subscriber Interface Enhancements

Introduced in LISTSERV Maestro 2.0, the Subscriber Interface allowed subscribers to join LISTSERV Maestro Hosted Lists and update their own information. Now, with version 2.1, LISTSERV Maestro improves the Subscriber Interface by adding a powerful and flexible Customization wizard, by providing the data administrator with new subscription management options, and by offering a powerful tool for translating the LUI, HUB, and Subscriber Interface.

The Subscriber Interface Customization Wizard

The Customization wizard for the Subscriber Interface makes it possible for the data administrator to customize every page and email that a subscriber will see. This will make it easier for you to focus on your subscriber's “touch points” and convey your organization's message. This same customization can be applied to all of the pages related to a single dataset or you can customize each individual page. For example, text for some lists can be in English, and text for others can be in French, or the entire look and feel of the signup page for an individual list can be different from all other lists in the dataset.

Context-sensitive instructions guide you through the Customization wizard, describing each template, listing the placeholders used in each template, and showing you how each template relates to the others.

The Customization wizard maintains two versions of every template – a draft version and a production version. You are also able to test the draft version of the template, using varying conditions, before releasing it into production.

For a summary on customizing the Subscriber Interface, see the Data Administrator's Manual for LISTSERV Maestro 2.1.

For step-by-step instructions on customizing the Subscriber Interface, see the Interface Customization Manual for LISTSERV Maestro 2.1. (This manual is new in version 2.1.)

New Subscriber Options

There are two new options for subscribers:

- Hosted LISTSERV List (HLL) subscribers can now control their “nomail” and digest settings using the Subscriber Interface.
- Hosted Recipient List (HRL) subscribers can now suspend mailings on lists that they subscribe to. This is similar to the “nomail” setting in Hosted LISTSERV Lists.

For more information on these new options, see the Data Administrator's Manual for LISTSERV Maestro 2.1.

New Resource Translation Tool

The Resource Translation Tool (RTT) is a separate application that makes it easy to make global changes in English or to translate all of the Subscriber, LUI, and HUB pages into another language. The languages that are available for the LUI and HUB interfaces are English, French, German, Spanish, and Swedish. There are 27 languages available in the Subscriber Interface. The default language for the Subscriber Interface can be set for the whole site or by group.

Before you can start using the Translation Tool, the following basic requirements must be met:

- You need to have access to a server with an installation of LISTSERV Maestro.
- You need to have an installation of the L-Soft Resource Translation Tool.
- You need to have a Translation Tool that matches the LISTSERV Maestro version that you have installed.
- You need a good understanding of editing HTML code (and the meaning of the various HTML tags and how to escape reserved HTML characters).

The Translation Tool contains several components – one for the LISTSERV Maestro Administration Hub, one for the LISTSERV User Interface Main Application, and one for the LISTSERV User Interface Subscriber Pages.

The Translation Tool can also be used to import translations from previous versions of LISTSERV Maestro.

For more information on the Translation Tool, see the Interface Customization Manual for LISTSERV Maestro 2.1.

Recipient Data Management Enhancements

Introduced in LISTSERV Maestro 2.0, hosted datasets and lists gave you the ability to manage your subscribers' data within LISTSERV Maestro. Now, with version 2.1, several new and powerful data management features have been implemented. These new features include: automatic bounce processing, flexible dataset administration, and shared subscriber management and reporting.

Automatic Bounce Processing

Automatic bounce processing has always been available for Hosted LISTSERV Lists (HLLs). Now, with version 2.1, LISTSERV Maestro has added the ability for you to define bounce processing settings for Hosted Recipient Lists (HRLs) in all of your datasets. For more information on defining automatic bounce processing, see the Data Administrator's Manual for LISTSERV Maestro 2.1.

Flexible Dataset Administration

In LISTSERV Maestro 2.1, numerous enhancements have been made to make dataset administration easier to perform. These enhancements include:

- The ability to make modifications to a “live” dataset, such as adding, changing, or removing data fields. Once you have existing datasets in the Recipient Warehouse, their settings can be fully edited until there are lists or members added to them. After lists or members become part of a dataset, only limited editing capabilities are available. For more information, see the Data Administrator's Manual for LISTSERV Maestro 2.1.
- The ability to designate “optional” fields as “hidden” or “read-only” so that subscribers cannot change them. It is then up to the data administrator or designated team member(s) to populate these fields. For more information, see the Data Administrator's Manual for LISTSERV Maestro 2.1.
- The Upload wizard can now import data from files that are missing mandatory data fields. This is done by simply defining default values to the fields with missing data. For more information, see the Data Administrator's Manual for LISTSERV Maestro 2.1.
- LISTSERV Maestro can now keep subscriber activity change logs to assist with recovery in the event of a “crash”. This log keeps track of the following activities – subscribing, unsubscribing, joining, un-registering, and address changes of all subscribers and members of hosted lists and datasets. By default, the subscriber change log is deactivated. For more information, see Section 13.2 [Subscriber Activity Change Log](#).
- The name fields for Hosted LISTSERV Lists (HLLs) can now be marked as “optional”, and the Firstname<space>Lastname restriction has been removed. For more information, see the Data Administrator's Manual for LISTSERV Maestro 2.1.
- The default language for use on the Subscriber Pages can now be set at the dataset level. To set the language, simply go to the General screen of the Recipient Dataset Definition wizard. For more information, see the Data Administrator's Manual for LISTSERV Maestro 2.1.
- The number of entries for any single lookup table has been increased from 200 to 1000 items. For more information about lookup tables, see the Data Administrator's Manual for LISTSERV Maestro 2.1.

Shared Subscriber Management and Reporting

Team Collaboration is now available for all datasets, allowing the data administrator to designate certain rights to other team members. These rights include the ability to add, remove, or modify subscribers, plus the ability to browse through subscriber data for purposes of downloading data for reports, etc. For more information on how to designate these rights to specific team members, see the Data Administrator's Manual for LISTSERV Maestro 2.1.

In addition, data administrators can now define default Team Collaboration settings for all new datasets. For more information, see the User's Manual for LISTSERV Maestro 2.1.

Tracking Enhancements

Several enhancements have been made to improve your ability to track the behavior of the recipients of your email jobs.

Tracking Restrictions

The administrator can now control what tracking types each user or group is able to define while tracking a LISTSERV Maestro email job. For more information on setting up system-wide default tracking restrictions, see the Section 7.2 [Application Default Settings](#). For more information on setting up the default tracking restrictions for a particular user or group, see Section 9.2.3 [Editing Component Specific Settings for Single and Group Users](#).

Tracking Links in Drop-Ins

If drop-ins are enabled in the content of a message and if the content contains at least one user-defined drop-in (not a system drop-in), then you now have the ability to track the links within this drop-in content. For more information on tracking link within drop-in content, see the User's Manual for LISTSERV Maestro 2.1.

Administration Enhancements

Several enhancements have been made to improve the administrator's ability to manage and control accounts. These enhancements are described in the following sections.

LUI Admin Account

The LUI interface for the special "admin" account has been improved, making it easier for the administrator to deal with staffing changes. These enhancements include:

- The ability to display jobs or reports owned by a particular owner or group. For more information, see Section 10.3 [Changing Job and Report Ownership](#).
- The ability to list all jobs or reports that have been "orphaned" when that particular account owner was deleted. For more information, see Section 10.3 [Changing Job and Report Ownership](#).
- The ability to change the owner of more than one job or report at a time. For more information, see Section 10.3 [Changing Job and Report Ownership](#).
- The ability to delete more than one job or report at a time. For more information, see Section 10.3 [Changing Job and Report Ownership](#).
- The ability to archive more than one job at a time. For more information, see Section 10.1 [Archiving Delivered Jobs](#).
- The ability to import more than one job at a time. For more information, see Section 10.2 [Importing Archived Jobs](#).

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 automatically loaded by the system during startup. However, in certain situations,

the administrator may want to refresh or get new versions of the translation files without restarting. This is now possible with the simple click of a button. For more information, see Section 8.5 [Refreshing the Subscriber Page Translations](#).

Creating a Test-Bed Backup

In certain situations (i.e. testing an upgrade, beta testing, or training), 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.

For more information, see Section 12 [Using a LISTSERV Maestro Test-Bed Backup](#).

DomainKeys Signatures

LISTSERV Maestro now allows you to use DomainKeys signatures to authenticate that the message (sent for a specific email job) does 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 will 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 allows you to 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.

For information on the default DomainKeys settings, see Section 24 [Authenticating Message Origin with DomainKeys Signatures](#).

For information on using DomainKeys Signatures while defining the sender information of an email job, see the User's Manual for LISTSERV Maestro 2.1.

System Support Enhancements

The following database and operating system enhancements have been added to LISTSERV Maestro 2.1:

- Asian languages are now supported by LISTSERV Maestro. This includes all Korean, Chinese, and Japanese characters (charsets). For more information, see Section 23 [Using International Character Sets](#).
- It is now possible to define a MS SQL Server database using its instance name.
- IBM DB2 version 8.2 is now a supported DBMS for a system database or a recipient data source. A new database plugin has been added to access IBM DB2 V8.2 with a “thin” driver,

which does not require the installation of the native DB2 database client on the LUI server. For more information, see Section 5.1.1 [The IBM DB2 V8.2 Thin Driver Database Plugin](#).

- Oracle 10g is now a supported DBMS for a system database or a recipient data source. For more information, see Section 5 [Defining External Database Connections](#) and Section 5.1.5 [The Oracle 8i, 9i, and 10g Thin Driver Database Plugin](#).
- The LISTSERV Maestro server is now supported on Mac OS X 10.4 and other compatible versions.
- The Mozilla Firefox browser on Mac OS X is now supported as a LISTSERV Maestro client for accessing the LUI and HUB interfaces. For more information on client system requirements, see Section 1.3 [Client System Requirements](#).
- The Safari browser (version 2.0, build 412 and later) on Mac OS X is now supported as a LISTSERV Maestro client for accessing the LUI and HUB interfaces. For more information on client system requirements, see Section 1.3 [Client System Requirements](#).
- MS SQL Server 2005 is now supported with the free open-source jTDS driver. The jTDS driver support is also available with MS SQL Server 2000. For more information, see Section 5 [Defining External Database Connections](#).
- New or customized web pages (independent of LISTSERV Maestro) can now be served by the Tomcat Server. For more information, see Section 25 [Adding Content to the Tomcat Server](#).

Section 1 Introduction to LISTSERV® Maestro Administration

Designed specifically to work with LISTSERV® 14.4 (or later), LISTSERV Maestro 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 also relies on the existence of an installation of **LISTSERV 14.4** (or later). 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.

By default, LISTSERV Maestro 2.1 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 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`

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

Once the program has been installed, set a compatible browser to:

`http://Your_LISTSERV_Host/hub`. Enter a password to log in and access the program. The default administrator password after a fresh installation is “admin”.



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

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.

This section of the manual contains a brief overview of the HUB interface. Navigation and functional icons are outlined. References to other sections of the manual containing greater detail and step-by-step procedures are provided.

The opening screen of LISTSERV Maestro's Administrator interface contains various sets of functional and navigational icons. At the top right corner of the screen there are two small icons, **Log Out** and **Help**. There are additional navigational icons on other pages within the application.

Table 1 Navigational Icons



Home – Brings the user back to the first screen, the LISTSERV Maestro HUB home page.



Log out – Ends the LISTSERV Maestro session and closes the account.



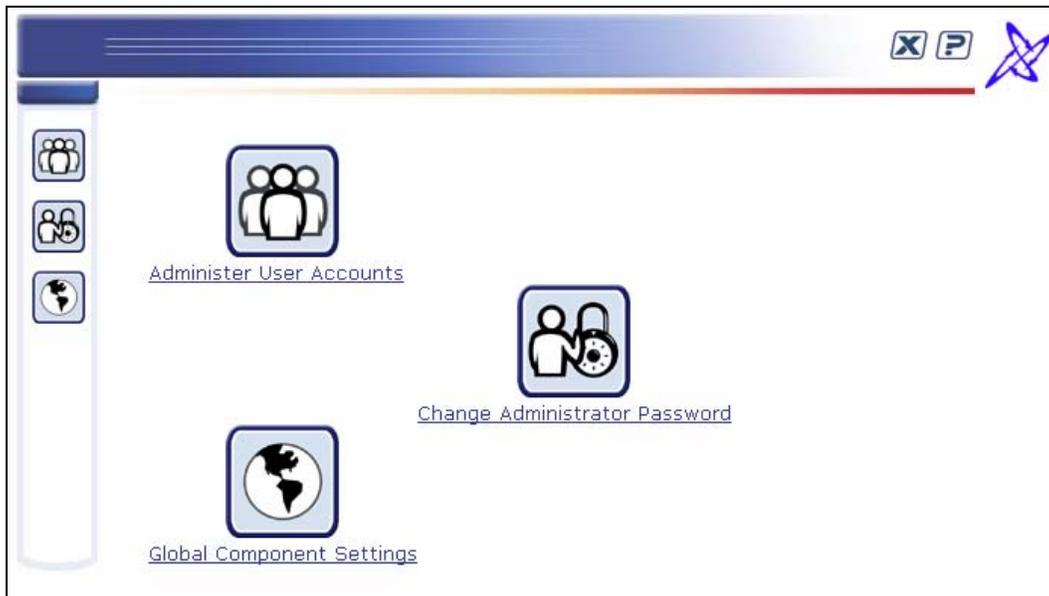
Up One Level – Brings the user up one level in the program, not necessarily back to the previous page.



Help – Provides access to page specific online help.

The center of the screen contains large icons that represent the major functions of the Administration Hub. These icons are also repeated along the left side of most screens within the program allowing for easy access from other parts of the administration interface.

Figure 1 Administration Hub Home Page



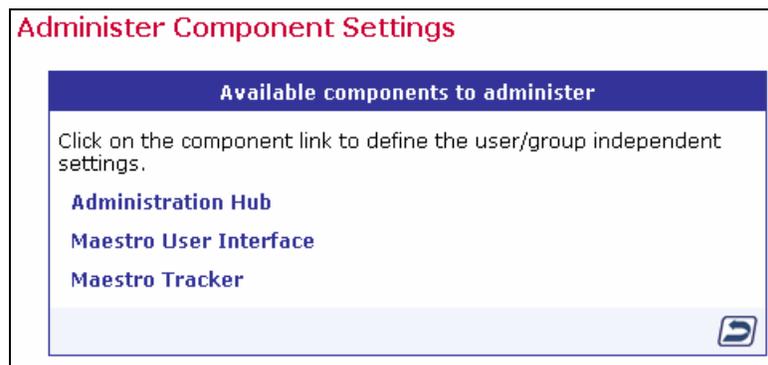
Administer User Accounts – Leads to the section of the application that allows the administrator to add new user accounts and groups, change user and group settings, and delete users and groups. User and group administration is documented in Section 9 [Creating and Administering User Accounts](#).

Change Administrator Password – Leads to the section of the application that allows the password used to log into the HUB and the special administrative user account to be changed. Changing the password is documented in Section 3 [Changing the Administration Password](#). The special administrative user account is documented in Section 10 [Special Administrative User Account](#).

Global Components Settings – Leads to the sections of the application that control the settings for each of the three components, the Administration Hub (HUB), the Maestro User Interface (LUI), and the Maestro Tracker (TRK).

Global component settings control the three components of LISTERV Maestro. Clicking on the **Global Component Settings** icon from the Home page opens a screen that is split into three main areas.

Figure 2 Global Component Settings



Administration Hub – Contains the settings for management of the HUB component.

- **Configure backup parameters for the HUB component, including:**
Setting the time for the daily backup.
Naming the backup folder.
Configuring any external processes to be run after backup.
Backup procedures are documented in Section 11 [Saving and Restoring a Backup](#).
- **Configure Administrative Email Notifications:**
Send or do not send email notification in the event of a system problem.
Send or do not send email notification for each system startup.
Assign an SMTP Host, SMTP Port, and sender address for notification email messages.
Assign email addresses to receive notifications.
Administrative Email Notifications are discussed in Section 8.5 [Administrative Email Notifications](#).

Maestro User Interface – Contains the settings for management of the LUI component. Settings entered here become the default settings for the system. Default settings can be overridden at the group and user levels. For more information on group and user levels, see Section 9 [Creating and Administering User Accounts](#).

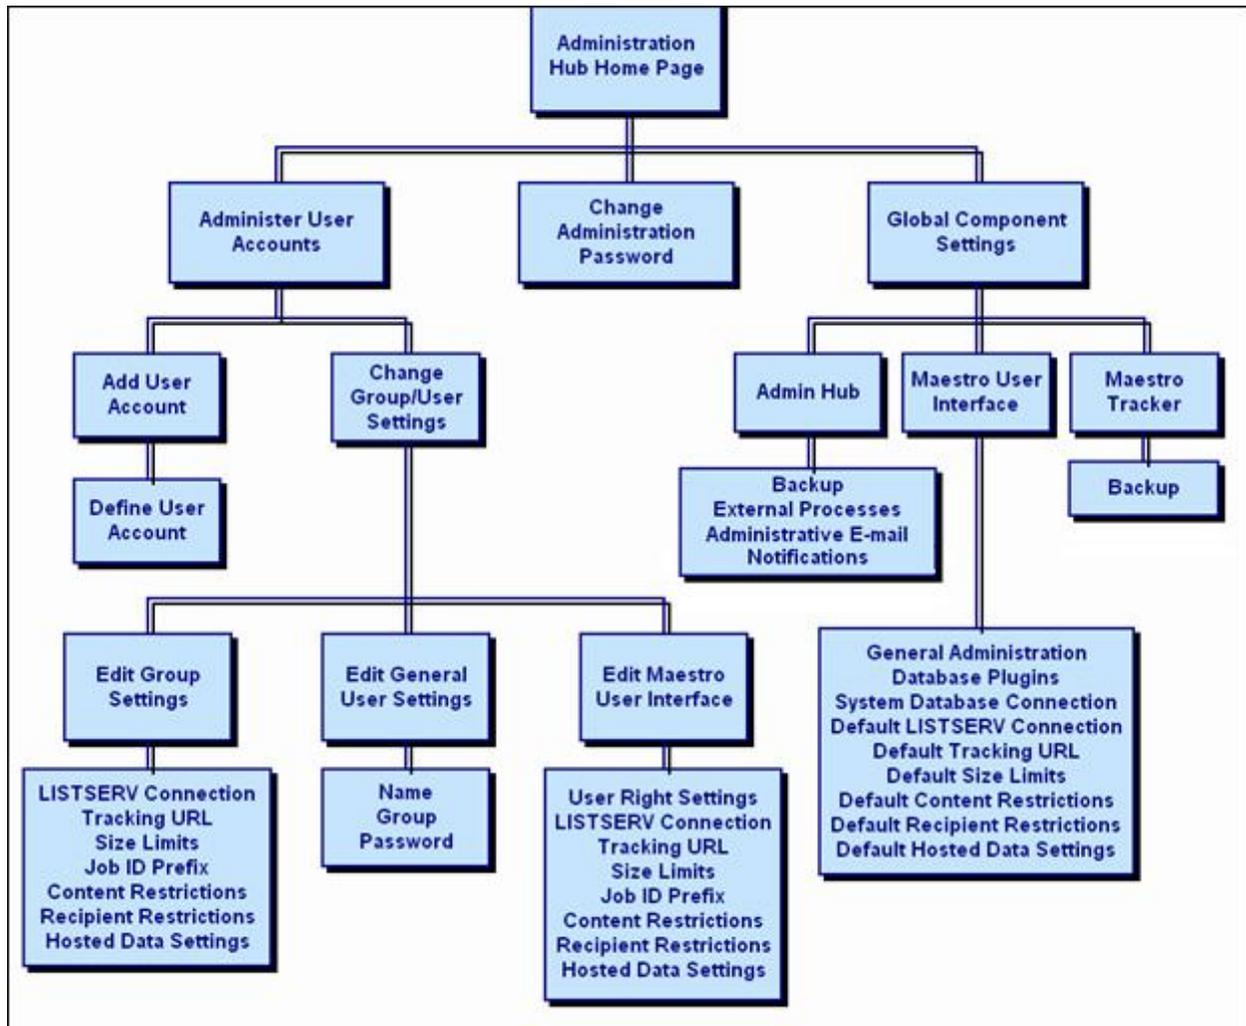
- **Configure General Settings for the LUI component, including:**
Naming the backup folder. Backup procedures are documented in Section 11 [Saving and Restoring a Backup](#).
Setting the event transfer interval from Maestro Tracker.
Naming the job archive folder.
Configuring Runtime Administration for restricting multiple logins, disabling sending, and locking the user interface. These settings are documented in Section 8.3 [Runtime Administration and System Shutdown](#).
- **Register Database Plugins:**
Adding a database plugin. Registering a database plugins is described in Section 5.2 [Registering a Database Plugin](#).
Deleting an existing plugin.
- **Configure the System Database Connection, including:**
Setting the maximum number of buffered connections.
Defining the system database as described in Section 4 [The System Database](#).
- **Define Default LISTSERV Connection:**
Setting the LISTSERV Host.
Setting the LISTSERV TCPGUI Port.
Entering the Client Address.
Entering the Client Password.
Setting the client address and password for Hosted LISTSERV Lists.
Enabling or disabling list archives for Hosted LISTSERV Lists.
Setting the folder location for list archives.
Viewing the database server name for Hosted LISTSERV Lists.
Setting up a dedicated bounce server.
More information on LISTSERV settings is documented in Section 7 [Settings for the Maestro User Interface](#).

- **Define the Default Tracking URL:**
Entering the name of the Tracker Host.
Setting the HTTP Port. The default tracking URL is documented in Section 7.3 [Setting the Default Tracker URL](#).
- **Set Default Size Limits:**
Setting the maximum size for a message.
Setting the maximum size for a file upload. For more information on size limits, see Section 7 [Settings for the Maestro User Interface](#).
- **Define Default Content Restrictions:**
Allowing or disallowing the use of AOL Rich Text formatting for the alternative part of an HTML message. For more information on AOL Rich Text, see Section 5.3.2 AOL Alternative Text for HTML Messages in the LISTSERV Maestro User's Manual.
Entering folders and/or URLs that the users are permitting to access for drop-in content elements. For more information on permitted folders and URLs, see [Section 7 Settings for the Maestro User Interface](#) or the online help.
- **Define Default Recipients Restrictions:**
Enabling, disabling, or hiding standard and advanced recipients types.
Entering folder name(s) that users are permitted to access for uploading "just-in-time" recipients lists. For more information, see Section 7 [Settings for the Maestro User Interface](#).
- **Define Default Tracking Restrictions:**
Enabling, disabling, and hiding tracking types.
For more information, see Section 7 [Settings for the Maestro User Interface](#).
- **Define Default Hosted Data Settings:**
Defining host name for subscriber access pages. Enabling or disabling the creation of Hosted Recipients Lists and Hosted LISTSERV Lists. For more information, see Section 7 [Settings for the Maestro User Interface](#).

Maestro Tracker – Contains the settings for management of the TRK component.

- **Configure backup settings:**
Setting the name of the backup folder.
Setting the number of previous backups to be kept. Backup procedures are documented in Section 11 [Saving and Restoring a Backup](#).
- **Configure the Communications Port:**
Section 14.1.3 [Ports used by Maestro Tracker](#) describes the ports used by Maestro Tracker.

Figure 3 Flow Chart of the Administration Hub



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. Execute at least the following configuration steps for security and access purposes:

1. **Change Administrator Password** – For security purposes, change the administrator password to something other than the default “admin” immediately after installation. Do not forget the new password; it is not recoverable. See Section 3 [Changing the Administrator Password](#) for more information.
2. **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 4 [The System Database](#).

3. **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 5 [Defining External Database Connections](#).

4. **Set up LISTSERV** – See Section 6 [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).

5. **Configure the Maestro User Interface (LUI)** – Two steps in particular need to be accomplished. These are also handles 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 7 [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 7.3 [Setting the Default Tracking URL](#) for information on how to define the default tracking URLs.
6. **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, whether any external processes will run before or after a backup, where the backup(s) will be saved.
 - **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.
 7. **Create User Accounts** – Create at least one user account that can be used to access the Maestro User Interface. See Section 9 [Creating and Administering User Accounts](#) for more details on how to proceed with creating user accounts, changing the password for the “sample” user account, or deleting the “sample” account.
 8. **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 18.3.1 [Securing the Trusted Root Certificate Keystore](#).

Section 3 Changing the Administrator Password

To change the administrator password, click the icon on the Administration Hub home page. 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 **[Save]** button to record the changes or the **[Cancel]** button to disregard.

Figure 4 Change Administrator Password Screen



Change Administrator Password

Old Password:

New Password:

Confirm Password:

OK Cancel

The default password after a fresh installation is "admin"

The default password after a fresh installation of LISTSERV Maestro is "admin".

Section 4 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
- **Oracle**
 - 8i
 - 9i
 - 10g

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).

4.1 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 4.2 [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 5.1 [Available Database Plugins](#) for details.
4. Register the appropriate "Plugin" in the Administration Hub (HUB) component. See Section 5.2 [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 Component Settings** icon. Next, click **Maestro User Interface**, and then click **System Database Connection**. The System Database Connection screen opens.

Figure 5 System Database Connection Screen

- 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.
- d. Fill out the database connection detail fields. These fields are different for each type of database selected.
- e. Click **[OK]** to submit the change.

Figure 6 System Database Connection Details Screen for DB2

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 4.4 [Removing or 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 4.4 [Removing or 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.

4.2 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.

4.2.1 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 5 [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 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.

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.

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.

4.2.2 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.

4.2.3 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.

4.2.4 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.

4.2.5 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
```

4.3 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.

4.4 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 (see Figure 5). 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 5 Defining External Database Connections

The Maestro User Interface component of LISTSERV Maestro uses a “system database” to store its working data, outlined in Section 4 [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 5.1 [Available Database Plugins](#) for details.
- Register the appropriate plugin in the Administration Hub (HUB) component. See Section 5.2 [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.

5.1 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 5.2 [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.

5.1.1 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
```

5.1.2 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 4.2.4 [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).

1. 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.

2. 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
```

5.1.3 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.

5.1.4 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.

5.1.5 The Oracle 8i, 9i, and 10g Thin Driver Database Plugin

These plugins are used for connecting to the Oracle database of version 8i, 9i, 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 8i 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.

5.1.6 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.

5.1.7 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.

5.1.8 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.

5.1.9 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.

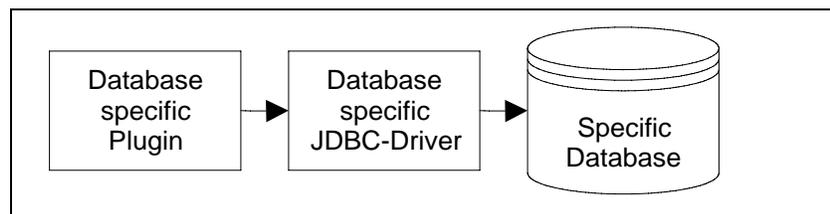
- **Plugin class name:** `com.lsoft.lui.db.odbc.ODBCDriverPlugin`
- **How to install the driver:** The driver is pre-installed together with LISTSERV Maestro.

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 5.2 [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 in Figure 7 below.

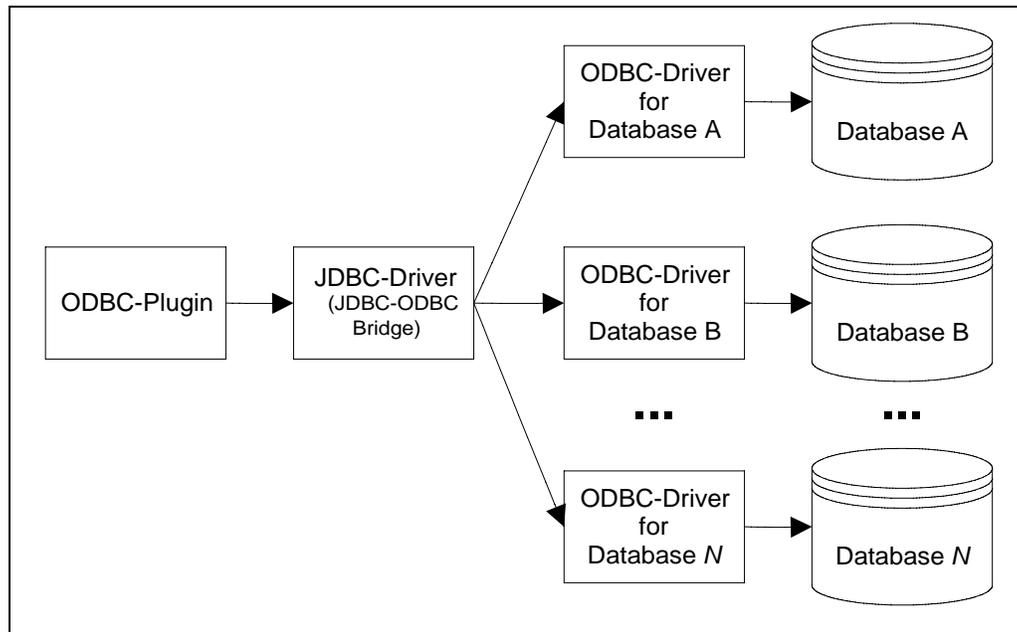
Figure 7 JDBC Driver Layers



In contrast, the ODBC driver plugin has one more layer shown in Figure 7. 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 8 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.

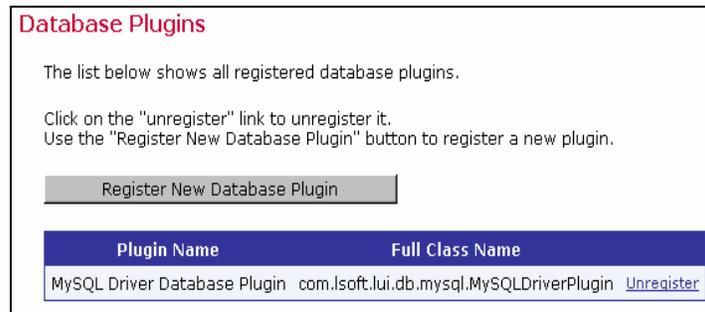
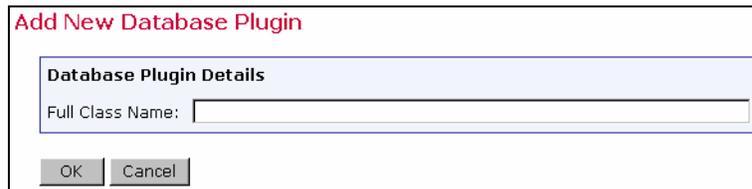


Tip: The term “database”, when used with ODBC, is interpreted quite broadly. ODBC drivers for data sources, such as plain text files or if Microsoft Excel files exist, turning them into “databases” in the sense that they can be used to create recipient lists and drop-in content. See the L-Soft White Paper entitled *Importing Data from Microsoft Excel into LISTSERV Maestro* for an example.

5.2 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 **Global Component Settings**, then **Maestro User Interface**, and then finally **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 5.1 [Available Database Plugins](#)).

Figure 9 Database Plugins*Figure 10 Add New Database Plugin*

Click **[OK]** 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 6 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 [9.2.3 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 also may 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 some additional configuration within LISTSERV in order to allow LISTSERV to connect to the user database and to allow LISTSERV Maestro to create new LISTSERV lists.

6.1 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 thus 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>.

6.1.1 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 7 [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.

6.1.2 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 2 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).
Linux Solaris HP-UX AIX	The external system database must be defined using Oracle or DB2.

If LISTSERV is running on this Operating System:	LISTSERV Maestro HLL support requirements are:
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 7 [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.

6.1.3 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 6.1.2](#)), 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\MySQL\bin\mysql.exe` and 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. View the database server name through the Maestro Administration Hub by clicking on Global Component Settings > Maestro User Interface > Default LISTSERV Connection. Under Connection Settings for Hosted LISTSERV Lists, look for the Database Server Name. This name should begin with `MAESTRODB_`, followed by a string of characters.

Figure 11 The Database Server Name

Connection Settings for Hosted LISTSERV Lists

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

Client Address for Hosted LISTSERV List Administration (Must be POSTMASTER):

Address: < default undefined >

Password: < default undefined >

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_dtn7nzn3**
(Must be configured as a database server name at the LISTSERV host)

In the example in Figure 11, the Database Server Name is `MAESTRODB_dtn7nzn3`. This server name uniquely identifies the Maestro system database to LISTSERV. Enter that database

server name into the LISTSERV 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 LISTSERV 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 LISTSERV 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 LISTSERV Maestro is being used, and LISTSERV 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 LISTSERV Maestro is being used and LISTSERV 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 much 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.

6.2 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 7 [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, see Figure 24). 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 (see Figure 34). 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 (see Figure 35). 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 (see Figure 33). 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 6.2.1 [Specifying a LISTSERV Host Name 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 6.2.2 [Specifying a Separate LISTSERV Instance for Processing Bounces](#).

6.2.1 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 pull down menu. An edit box will appear. Enter the host name in the edit box.

Figure 12 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: papaya.ease.lsoft.com

External Host Name:
 Default: papaya.ease.lsoft.com

LISTSERV TCPGUI Port: Default: 2306

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

Address: Default: maestro@papaya.ease.lsoft.com

Password: Default: *****

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.

6.2.2 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 (see Figure 18).

If the dedicated bounce processing host has only a single name, enter that name into the **LISTSERV Host** field of the LISTSERV Connection screen (see Figure 18). 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 6.2 [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 (see Figure 18).

6.3 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 7 Settings for the Maestro User Interface

To select settings for the Maestro User Interface, click on the **Global Component Setting** icon from the home page of the Administration User Interface. Next, click **Maestro User Interface**.

Figure 13 Global Component Settings for Maestro User Interface



The top three selections define application settings, and the bottom six selections define application defaults. Click on a link to define or edit the settings.

7.1 Application Settings

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

General Administration – Defines general Maestro User Interface settings.

Figure 14 General Administration of the Maestro User Interface

General Administration of Maestro User Interface

General Settings

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

Event transfer interval (from Maestro Tracker): min.

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

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:

- **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.
- **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 text message at the top of each screen to actively 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.

Database Plugins – Register and unregister database plugins. Database plugins allow LISTSERV Maestro to communicate with databases. For more information, see Section 5.2 [Registering a Database Plugin](#).

Figure 15 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.

Register New Database Plugin

Plugin Name	Full Class Name
MySQL Driver Database Plugin	com.lsoft.lui.db.mysql.MySQLDriverPlugin Unregister

System Database Connection – Defines the settings for the Maestro System Database. For more information on the system database, see Section 4 [The System Database](#).

Figure 16 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:

MySQL User Name:

Password:

Database Host Name:

TCP/IP Port:

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

OK Cancel

- **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.

7.2 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.

Default LISTSERV Connection – 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.



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 6.2.1 [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 6.2.1 [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.
- **Client Address for Mail Delivery** – Enter the email address that has been configured in LISTSERV to have the right to send DISTRIBUTE jobs. See Section 6.1 [Preparing LISTSERV to Process Jobs from LISTSERV Maestro](#).
- **Client Password for Mail Delivery** – Enter the password configured with the LISTSERV Client Address. See Section 6.1 [Preparing LISTSERV to Process Jobs from LISTSERV Maestro](#).

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

- **Client Address for Hosted LISTSERV List Administration** – This should be the email address that was added as a `POSTMASTER` in the LISTSERV Site

Configuration file (see section 6.1.2 [Preparing LISTSERV to Allow List Creation by LISTSERV Maestro](#)).

- **Password** – Enter the LISTSERV password that was registered for the Client Address entered above.
- **List Archive Settings for LISTSERV Hosted 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 6.1.3 [Preparing the LISTSERV List Archive Folder](#).

Figure 17 Default LISTSERV Connection

LISTSERV Connection

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: ▾

LISTSERV TCPGUI Port: Default: 2306

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

Address: <default undefined>

Password: <default undefined>

Connection Settings for Hosted LISTSERV Lists

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

Client Address for Hosted LISTSERV List Administration (Must be POSTMASTER):

Address: <default undefined>

Password: <default undefined>

Usage of List-Archives for Hosted LISTSERV Lists:

▾

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

Database Server Name: **MAESTRODB_1Weto8**
(Must be configured as a database server name at the LISTSERV host)

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)

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 6 [LISTSERV and LISTSERV Maestro](#).

Figure 18 Dedicated Bounce Server

LISTSERV Connection for Automatic Bounce Handling

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

Inherit settings from application defaults
 (Currently: Use standard LISTSERV server, see above)

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

Use dedicated server as follows:

LISTSERV Host: < default undefined >

External Host Name: < default undefined >

LISTSERV TCPGUI Port: Default: 2306

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

Address: < default undefined >

Password: < default undefined >

Default Tracking URL – Generates the tracking URL for all accounts where no explicit tracking URL is defined on either the single user or group level. For more information on the default tracking URL, see Section 7.3 [Setting the Default Tracking URL](#).

- **Tracker Host** – Enter the host name of the server running the Maestro Tracker component.
- **HTTP Port** – Enter the port where the Maestro Tracker component on the Maestro Tracker host listens for HTTP connections. The default port number is 80.

Figure 19 Default Tracking URL

Tracking URL

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

Tracker Host: < default undefined >

HTTP Port: Default: 80



Important: All accounts and groups must use tracking URLs that point to the same physical Maestro Tracker server, using the same HTTP port. Although it is possible to enter different Maestro Tracker host names and port settings on the same global application, group level, or individual user level, all those entries must point back to the same physical server, using the same HTTP port. For more information, see Section 7.3.1 [Multiple Tracking URLs](#) and the online help.

Default Size Limits – 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 20 Default Size Limits

Default Content Restrictions – Define 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 more information on using this setting.

Figure 21 Default Content Restrictions

Default Recipients Restrictions – The 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. 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.

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-in-time" CSV files for recipients definitions. If left blank, CSV files used for just-in-time recipients definitions in the recipients definition wizard will not be allowed.

Figure 22 Default Recipients 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.

Standard recipients types
 Enable at least one of the following standard recipients types.

	Enabled	Disabled	Hidden
Use Existing Recipients Target Group	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Send to an Existing LISTSERV List	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Upload a Recipients Text File	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Select Recipients From a Database	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Let LISTSERV Select Recipients From a Database	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

Advanced recipients type
 You can additionally enable the following advanced recipients type.

	Enabled	Disabled	Hidden
Determine Recipients by Inspecting the Reaction on another Job	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

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

Default Tracking Restrictions – 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**.

Figure 23 Default Tracking Restrictions

Tracking 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"/>

OK Cancel

Default Hosted Data Settings – Defines the settings for hosted recipient data.

- **General Hosted Data Settings** – Defines an external host name under which subscriber access pages are accessible to subscribers. This should be the same host that the Maestro User Interface (LUI) is running on, but if the host has a different external host name, enter it here.
- **Hosted List Settings** – Select whether or not users/groups may create Hosted Recipient Lists or Hosted LISTSERV Lists.

Figure 24 Default Hosted Data Settings

Hosted Data Settings

General Hosted Data Settings

Define the host name and HTTP-port under which the subscriber access pages of the datasets are accessible to subscribers (leave the field empty to use the default):

External Host Name: Default: Name of Maestro User Interface host and port 80

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.)

OK Cancel

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.

Default DomainKeys Settings – Defines whether or not to use DomainKeys signatures to authenticate the origin of the LISTSERV Maestro email messages. See Section 24 [Authenticating Message Origin with DomainKeys Signatures](#).

7.3 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 host name and port that LISTSERV Maestro will use for the tracking URLs generated for the message. 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&
```

All account holders who do not have separate single user or group tracking URLs configured use the default tracking URL. The administrator can define the tracking URL host and port on an 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).

To define the tracking URL host and HTTP-port:

- **Default Tracking URL** – For all accounts and groups that do not have individual settings, click **Global Component Settings > Maestro User Interface > Default Tracking URL**. The settings defined here will affect all accounts that do not have a tracker host or HTTP-port configured on a single user or group level.
- **Single User Tracking URL** – For an individual account that is not part of a group, click **Administer User Accounts**. Click on the user name of the account to be

configured (must be an account without a group), and then click on **Maestro User Interface**. The settings defined here affect only the selected account.

- **Group Tracking URL** – For all accounts in a group, click **Administer User Accounts**, and then click on the group name of any account that is a member of the group to be configured. The settings defined here affect all accounts in the selected group.

For all three choices, there are two possible settings:

- **Tracker Host** – Enter the host name to be used in the tracking URL. Leave empty to accept the default (if any). If left empty with no default given, users of this account or group will not be able to send email messages with tracking.
- **HTTP Port** – Enter the HTTP port to be used in the tracking URL. Leave empty unless using a proxy (see the next section for more information).

7.3.1 Multiple Tracking URLs

There are many uses for setting up multiple tracking URLs. One use of multiple tracking URLs is to customize the look of the tracked URL within the message. Another use of separate tracking URLs would be in the case of using a proxy where the host name of the tracker host used in the internal network is different from the host name that external users would have to use.

Customization is useful in an environment where the same LISTSERV Maestro setup is shared between several distinct groups that want to differentiate themselves in the tracking URL that is merged into the messages 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`. These group accounts are created in order to be used by the members of the two divisions (team collaboration can also be employed 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 messages that are delivered. Therefore, the administrator also sets up two DNS names, `electro.mycorp.com` and `homeapp.mycorp.com` as aliases for the same server. Next, the administrator uses the procedure described above to set the tracker host name of the `electro` and `homeapp` groups to the corresponding aliases. As a result, even though both divisions are sharing the same LISTSERV Maestro installation, the tracking URLs generated for the email they send are “customized” to contain a host name that matches each division’s name.

Another example is 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 administrator would set the proxy’s host name and the proxy-port that is forwarded to the HTTP port on the tracker host as the default “Tracker Host” and “HTTP Port” setting, to be used by all accounts.



Note: HTTP Port should always be left empty unless there is a proxy at the new host/port pair that redirects the connection to the single port on the TRK server that processes all tracking events, as described in Section 7.3.1 [Multiple Tracking URLs](#).

Section 8 Administrative Policies

Every institution and business using networked 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 8.4 [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.

8.1 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.

8.2 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 located under **Global Component Settings > Maestro User Interface > General Administration**. Use these settings to reach a safe shutdown state, where shutting the system down will not disrupt jobs being sent or users in the midst of preparing a job. Follow these steps to shut down LISTSERV Maestro safely:

1. Open a web browser and access the Maestro User Interface, logging in with the administrative account. See Section 10 [Special Administrative User Account](#) for more information.
2. Open a second browser and access the Administration Hub. Go to **Global Component 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, click on the **Outbox** icon. 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.

8.3 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.

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 9.2 [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 9.2 [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 9.2 [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 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-in-time" CSV files for recipients definitions. If left blank, CSV files used for just-in-time recipients definitions in the recipients definition wizard will not be allowed.

- **Hosted List Restrictions** – Under **Global Component 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 6 [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 Component 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.

8.4 Administrative Email Notifications

A link to configure administrative email messages is located on the Component Administration for the Administration Hub screen.

Figure 25 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 be sent to each configured recipient address. To have administrative email notifications sent, select the option to send email notifications. If notification is desired for system start, check **Send a notification email at each startup**. The following settings need to be configured to use email notifications:

- **SMTP Host** – Enter the host name for the SMTP server that will be used to mail the notifications. This is a mandatory field and must be filled out with a valid host name. This host name will be used to send mail by the LISTSERV Maestro component that encounters the problem. Make sure to specify a host name that is reachable from all servers running a LISTSERV Maestro component.
- **SMTP Port** – Enter the SMTP port on which the SMTP server on the host specified listens for SMTP connections. This field is optional. If left empty, the standard SMTP port 25 is used.

- **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.

8.4.1 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.

8.5 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. Go to **Global Component 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.

Figure 26 Refreshing the Subscriber Page Translations

The screenshot shows the 'General Administration of Maestro User Interface' window. It is divided into several sections:

- General Settings:** Includes fields for 'Backup folder:' (with a note '(Leave empty for default: "backup".)'), 'Event transfer interval (from Maestro Tracker):' (set to 10 min), and 'Job archive folder:' (with a note '(Leave empty for default: "archive".)').
- Runtime Administration:** Includes checkboxes for 'Multiple Logins' (Disallow multiple logins with the same user account.), 'Outbox' (Sending is disabled.), and 'Login Access' (LISTSERV Maestro User Interface is locked.). Below these are two text boxes for messages shown during login lockouts.
- Subscriber Access Page Translations:** This section is highlighted with a red box. It contains the text: 'Subscriber access page translations are automatically retrieved by the system during startup. Click the button below to refresh the page translations now, during runtime.' and a button labeled 'Refresh Translations Now'.

At the bottom of the window are 'OK' and 'Cancel' buttons.

For more information on the Resource Translation Tool, see the Interface Customization Manual for LISTSERV Maestro 2.1.

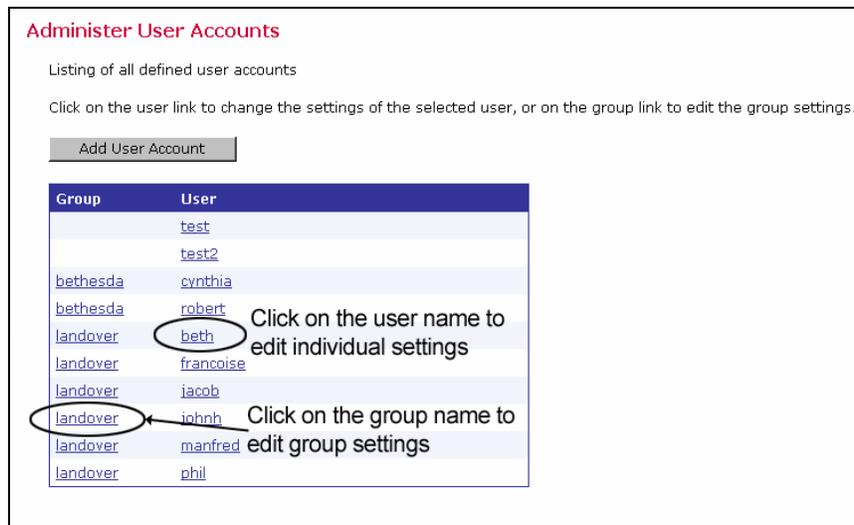
Section 9 Creating and Administering User Accounts

The Administer User Accounts screen displays a list of all defined user accounts. Each user name is a link that opens the Define User Account screen. Links on this 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 Administer User Accounts screen it is possible to:

- Add a new user account
- Assign users to a group
- Edit existing user account settings, including deleting the accounts
- Edit existing group settings

Figure 27 Administer User Accounts



9.1 Creating a New User Account

To create a new user account, click the **[Add User Account]** button.

- **User** – 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.

- **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.

Figure 28 Defining User Account

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.

9.2 Editing Account Information and Assigning Single User Settings

Once a new account has been created, it must be edited to assign single user settings. To edit existing account information, click on the account you want to edit. The Account Administration screen opens. Click on **User name, group and password** to edit a user's name, group, and password. Alternatively, select **Maestro User Interface** to control how the user interacts with the Maestro User Interface.

Figure 29 Editing Account Information

9.2.2 Editing General User Settings

To change the user name, group, or password of an existing account, click on **User name, group and password**. The Define User Account screen opens. To keep an old password while changing other settings, leave both password fields blank. To delete an account, click the **[Delete]** button.

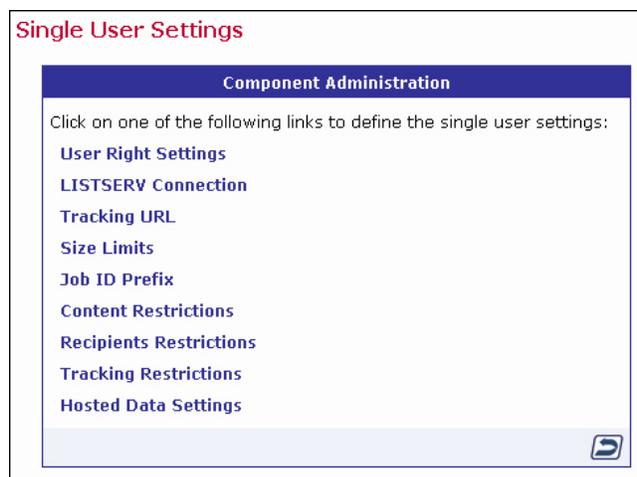
9.2.3 Editing Component Specific Settings for Single and Group Users

After an account is initially created, click on **Maestro User Interface** to open the selection list for all user specific settings. This list will vary depending on whether the user is a member of a group (a group user) or not (a single user). The selection list displayed in Figure 30 is for a single user who is not part of a group. The selection list displayed in Figure 31 is for a group user.

To edit settings at the group level, click on the name of the group from the Administer User Accounts screen (see Figure 277).

Click on any setting to open a screen to edit the setting.

*Figure 30
User Settings for a Single (Non-Group) Member*



*Figure 31
User Settings for a Group Member*



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 members in a group.
- **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 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.

New users will default to having all privileges granted except the right to **administer Target Groups** and **administer hosted Recipient Datasets**. (For more information on target groups and hosted recipient datasets, see the LISTSERV Maestro Database Administrator's Manual.) If certain privileges are revoked, the user will see a slash through the icon(s) in the Maestro User Interface that accesses the corresponding function(s).

Figure 32 User Right Settings

LISTSERV Connection – The LISTSERV Connection can be set 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 7.2 [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 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 **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 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 set 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 have all LISTSERV Connection settings defined (see Figure 33). Group user accounts can only define the **LISTSERV Client Address** and **LISTSERV Client Password** settings (see Figure 344). 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.

Figure 33 Default LISTSERV Connection for a Group with Dedicated Bounce Server

Currently selected group: landover (Settings for component: LUI)

LISTSERV Connection

Use the information on this line to determine which settings are being edited. Group settings become the default for the group, overriding Global default settings.

LISTSERV Connection

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

LISTSERV Host: Default: papaya.ease.isoft.com

External Host Name: ▼

LISTSERV TCPGUI Port: Default: 2306

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

Address: Default: maestro@papaya.ease.isoft.com

Password: Default: *****

Connection Settings for Hosted LISTSERV Lists

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

Client Address for Hosted LISTSERV List Administration (Must be POSTMASTER):

Address: Default: maestro-admin@papaya.ease.isoft.com

Password: Default: *****

Usage of List-Archives for Hosted LISTSERV Lists:

▼

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

Database Server Name: **MAESTRODB_dtn7nzn3**
(Must be configured as a database server name at the LISTSERV host)

LISTSERV Connection for Automatic Bounce Handling

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

Inherit settings from application defaults
(Currently: Use standard LISTSERV server, see above)

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

Use dedicated server as follows:

LISTSERV Host: <default undefined>

External Host Name: ▼

LISTSERV TCPGUI Port: Default: 2306

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

Address: <default undefined>

Password: <default undefined>

Figure 34 LISTSERV Connection for a Single User without a Dedicated Bounce Server

Currently selected user: sample (Settings for component: LUI)

LISTSERV Connection Use the information on this line to determine which settings are being edited.

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

LISTSERV Host: Default: papaya.ease.lsoft.com

External Host Name: ▾

LISTSERV TCPGUI Port: Default: 2306

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

Address: Default: maestro@papaya.ease.lsoft.com

Password: Default: *****

Connection Settings for Hosted LISTSERV Lists
 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.
 Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

Client Address for Hosted LISTSERV List Administration (Must be POSTMASTER):

Address: Default: maestro-admin@papaya.ease.lsoft.com

Password: Default: *****

Usage of List-Archives for Hosted LISTSERV Lists:

▾

Inherited base folder for archive location: e:\LISTSERV\LISTS

Database Server Name: **MAESTRODB_dtn7nzn3**
 (Must be configured as a database server name at the LISTSERV host)

LISTSERV Connection for Automatic Bounce Handling
 Leave the fields empty to use the defaults, if defined, or enter a value to overwrite the default.

Inherit settings from application defaults
 (Currently: Use standard LISTSERV server, see above)

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

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

OK Cancel

Figure 35 LISTSERV Connection for a Group User

Currently selected user: liam (Settings for component: LUI)

LISTSERV Connection

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

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

Address: Default: maestro@papaya.ease.lsoft.com

Password: Default: *****

OK Cancel

For information and instructions on how to fill out the fields for setting the LISTSERV Connection, see Section 6 [LISTSERV and LISTSERV Maestro](#). For information about setting a special external host name, see Section 6.2.1 [Specifying the LISTSERV Host with Different Internal and External Names](#). For information on setting up a dedicated LISTSERV instance for processing bounces, see Section 6.2.2 [Specifying a Separate LISTSERV Instance for Processing Bounces](#).



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

Tracking URL – Enter the Tracker URL for the user. Each user or group can have a different Tracker URL, although they all must lead back to the same tracker component. The domain name must resolve to an IP address on the server where the tracker component is installed.

Figure 36 Tracking URL

Size Limits – Sets a size limit for email messages. 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.

Figure 37 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 chargeback 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 38 Job ID Prefix

Job ID Prefix

Job ID Prefix Settings

Never use a job ID prefix
 Any job ID prefix allowed
 Use the following prefix for every job ID:
 Only job ID prefixes from the following list are permitted:

temp
news
sale

Enter one prefix per row.

OK Cancel

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.

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 39 Drop-In 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
 Use the inherited setting: 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

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.

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 40 Recipients Restrictions

Recipients Restrictions

Recipients Type Restrictions

Define which recipients types are available to the user.
 Select "Use Default" to inherit the default settings from application level.
 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.

Standard recipients types
 Enable at least one of the following standard recipients types.

	Enabled	Disabled	Hidden	Use Default	
Upload a Recipients Text File	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	Default: Enabled
Use Existing Recipients Target Group	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	Default: Enabled
Select Recipients From a Database	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	Default: Enabled
Send to an Existing LISTSERV List	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	Default: Enabled
Let LISTSERV Select Recipients From a Database	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	Default: Enabled

Advanced recipients type
 You can additionally enable the following advanced recipients type.

	Enabled	Disabled	Hidden	Use Default	
Determine Recipients by Inspecting the Reaction on another Job	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	Default: Enabled

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

To set component settings for a group, click on the group name. The Group Settings screen opens and will list the component settings:

- Default LISTSERV Connection
- Tracking URL
- Size Limits
- Job ID Prefix

- Content Restrictions
- Recipient Restrictions
- Hosted Data Settings
- Tracking Restrictions

The screens that configure these settings are very similar to the screens that configure single user settings. The 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.

A variation to this rule is the group default LISTSERV Connection. 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). To set 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 have all LISTSERV Connection settings defined (see Figure 33). Group user accounts can only define the **LISTSERV Client Address** and **LISTSERV Client Password** (see Figure 34).

Tracking Restrictions – This option 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 use the default setting, select **Use Default**.

Figure 41 Tracking Restrictions

Tracking Restrictions

Define which types of tracking are available to the group members.
 Select "Use Default" to inherit the default settings from application level.
 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	Use Default	
Personal Tracking	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	Default: Enabled
Anonymous Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	Default: Enabled
Unique Tracking	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	Default: Enabled
Blind Tracking	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	Default: Enabled

OK Cancel



Note: For information on setting the default tracking restrictions, see Section 7.2 [Application Default Settings](#).

Hosted Data Settings – Defines the settings for hosted recipient data.

- **General Hosted Data Settings** – Defines an external host name under which subscriber access pages are accessible to subscribers. This should be the same host that the Maestro User Interface (LUI) is running on, but if the host has a different external host name, enter it here. If left blank, then the default will be used.

- **Hosted List Settings** – Select the type of hosted list the group members can create.

Figure 42 Hosted Data Settings

Hosted Data Settings

General Hosted Data Settings

Define the host name and HTTP-port under which the subscriber access pages of the datasets are accessible to subscribers (leave the field empty to use the default):

External Host Name: Default: Name of Maestro User Interface host and port 80

Hosted List Settings

Define which types of Hosted Lists may be created by the members of the selected group:

Use inherited Application Defaults:
The group members may create Hosted Recipient Lists and Hosted LISTSERV Lists.

Define Group Specific Settings:

The group members may create Hosted Recipient Lists

The group members may create Hosted LISTSERV Lists

Hosted LISTSERV lists are created with the name suffix specified below:

Suffix: (If no suffix is specified, the users in the group will not be allowed to create hosted LISTSERV lists.)

Note: If Hosted LISTSERV Lists are enabled for this group, then the settings in the sub-section "Connection Settings for Hosted LISTSERV Lists" of the LISTSERV Connection must also be configured.

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 7.2 [Application Default Settings](#).

9.3 Deleting a User Account

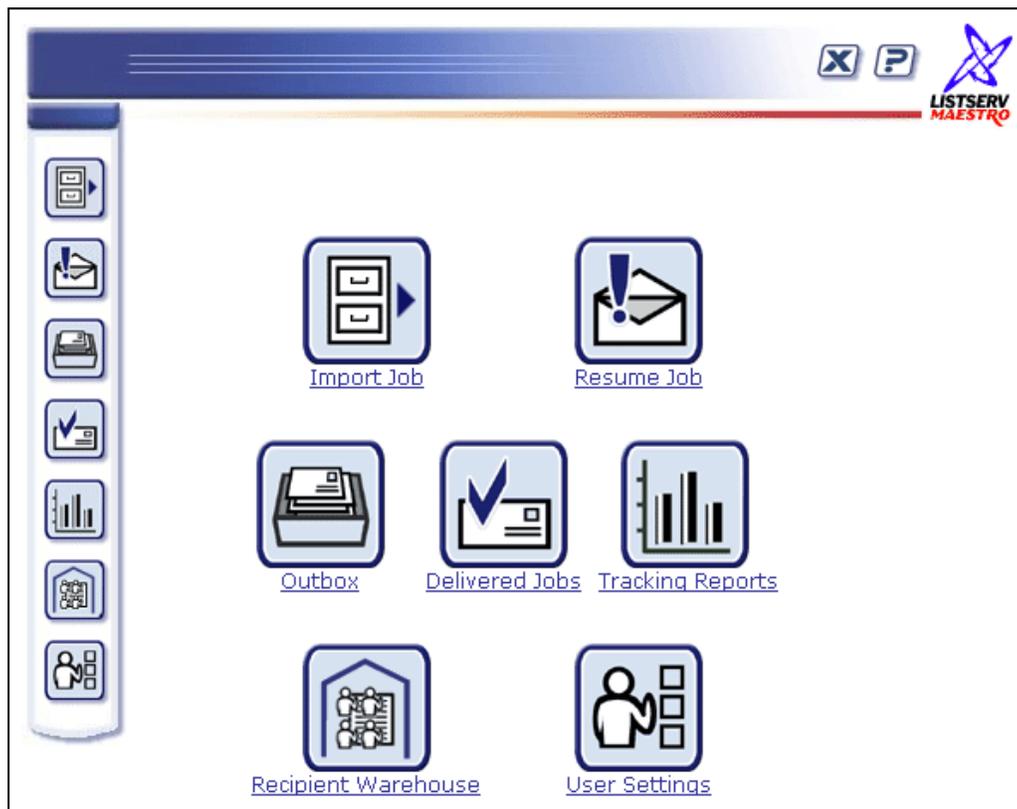
To delete a user account, select the user from the Administer User Accounts screen. Click on **User name, group and password**, and then click on the **[Delete]** button. Confirmation of the selection prior to the deletion of the account is required. Any jobs, reports, sender profiles, and drop-in content elements owned by that user become “orphaned” and need to be reassigned. See section 10, [Special Administrative User Account](#).

Section 10 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, log into the Maestro User Interface (LUI) as an administrator, using `admin` in the **User** field, and then typing in the administrator password that was configured in the Administrator Hub. The home page for the administrative user account is different from that of normal users.

Figure 43 Special Administrative User Account



10.1 Archiving Delivered 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 under the **Import Job** icon 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 7 [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 the **Delivered Jobs** icon from the home page of the administrative user account. The Delivered 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 Delivered Jobs screen.

– OR –

- Check the box next to each job ID that you want to archive, and then click the **Actions** link. Select **Move all selected jobs to archive**. The Job Administration screen opens. Click the **[Move to Archive]** button. Click the **[Cancel]** button to cancel the operation and return to the Delivered Jobs screen. By using this method of archiving, you can archive more than one job at a time.

Figure 44 Archiving a Job from the Delivered Jobs Screen

Delivered Jobs

All previously delivered jobs for your account.

Display the jobs owned by: ▼

Jobs that:

- ✓ were processed successfully
- ✗ were processed previously and failed during delivery

Click on the Job ID to open a job or choose from the actions menu to manipulate several jobs at once.

[Actions](#)

Job ID	Job Title	Mail Type	Date/Time of Delivery	State	
<input type="checkbox"/> 051012A	test	Plain	Oct. 13, 2005 11:58:56 AM	✓	Delete

10.2 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 on the **Import Job** icon located on the home page. A listing of all the jobs currently present in the archive will appear. Then, 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 Delivered Jobs 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 the **Actions** link. Select **Import all selected 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 Delivered Jobs screen, check the box next to the job(s) you want to assign a new owner to, and then click the **Actions** link. Select **Change owner of all selected jobs**, and then use the drop-down menu to select an owner for the imported archived job. Once restored, an imported job will be listed again in Delivered Jobs 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 45 Importing a Job from the Archived Jobs Screen



10.3 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 in the Outbox
- 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), 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, click on the appropriate icon to open the listing for that job. Click **Resume Job** for a listing of all the current jobs that have not been delivered yet, **Outbox** for all the jobs that have been authorized for delivery, or **Delivered Jobs** for all the jobs that have already been delivered. To reassign a report, click the **Tracking Reports** icon for a listing of all the currently defined reports.

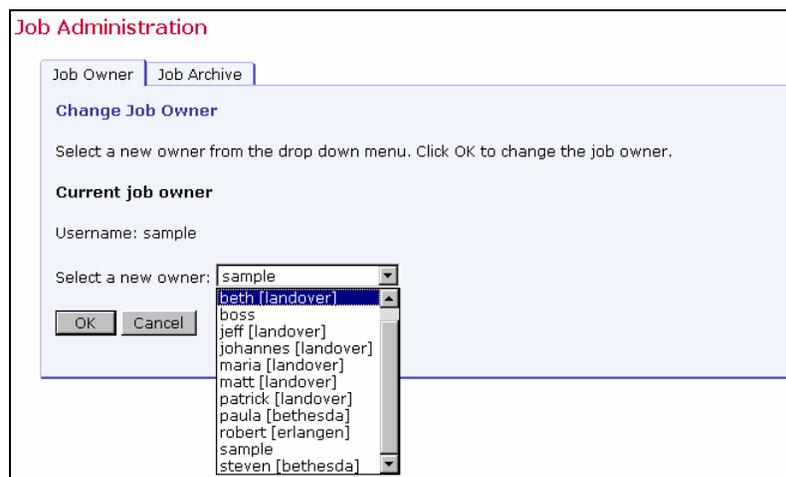
If you’ve clicked on the **Outbox** or the **Delivered Jobs** icon, then a listing of all the jobs currently present for that icon will appear. If you want to view a job owned by a specific user, click the drop-down menu and select that particular owner. Then, you can:

- 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 the **Actions** link. 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 46 Change Job Owner from Job Owner Tab



If you've clicked on the **Tracking Reports** icon, then a listing of all the reports currently present for that icon will appear. 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.

Figure 47 Change Report Owner Screen

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.

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.

10.4 Changing Sender Profile and Drop-In Content Element Ownership

Drop-in content elements and sender profiles are utility items in LISTSERV Maestro created for the convenience of the users. These items, if created by a user who is not part of a group, are owned by that single user and cannot be used by any one else. If a member of a group creates items, then everyone in that group can use them. In addition, users in that group who have the necessary right can create new items, and delete or modify existing items.

An item can lose its owner and become ownerless under certain circumstances:

- If an item is owned by a non-group account, and that account is deleted.
- If a group owns an item, and the last account of that group is deleted or is moved out of the group (into another group or made into an account without a group).

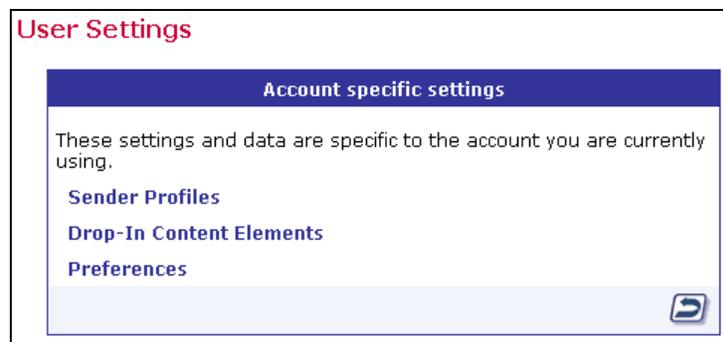
The ownership of any item, whether it already has an owner or not, can be changed by the administrator using the special Administrator User Account. The administrator can also delete any item.

To change the ownership of a Sender Profile or Drop-in Content Element, click on the **User Settings** icon from the homepage. The User Settings screen opens. Select the item type to change from the menu.



Note: The **Preferences** setting is used to set the user preferences for the Administrator User Account, and has nothing to do with changing the settings or item ownership for other accounts.

Figure 48 Change Ownership



Clicking on any of the menu entries will open a list of the related items, grouped after their current owners. At the top of the list, a drop-down list with all available owners appears. The list includes all available non-group accounts, all available groups, and a special <No Owner> entry. Below the drop-down list is a table listing all items that are currently owned by the selected owner account or group.

Figure 49 Select Account/Group and Item to Change Ownership

Manage Drop-In Content Elements

All defined drop-in content elements for the selected account/group:

Group demo

Click on a drop-in element to change the owner or to delete it.

Name	Description	Date/Time saved
new_address_html	New Company Address - html	May 21, 2004 06:42:44 PM
phone	company phone number	May 21, 2004 06:43:15 PM
recall_html	Address to send recalled products	May 21, 2004 06:41:13 PM

Click on an item to open a screen that lets you select a new owner or delete an item.

Figure 50 Change Drop-In Owner

Change Drop-In Owner

Selected Drop-In: new_address_html (New Company Address - html)

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

Current drop-in owner

Group: demo

Select a new owner for the drop-in...

Select a new owner: Group demo

Account sample
Group brandnew
Group demo
Group landover
Group marketing
Group newdemo
Group sales
Group senate-test
Group tc

OK Cancel

Delete

...or delete the drop-in.

If an item is moved to a new owner, and the new owner already has an item with the same name, 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.

10.5 Changing Recipient Dataset and Target Group Ownership

As with Sender Profiles and Drop-In Content Elements, Recipient Datasets and Target Groups may have either single user or group ownership. Single user Datasets and 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 ownerless under certain circumstances:

- If an item is owned by a non-group account, and that account is deleted.
- If a group owns an item, and the last account of that group is deleted or 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 Recipient Dataset or Target Group, click on the **Recipient Warehouse** icon from the homepage. The Recipient Warehouse screen opens. Select the item type to change from the menu.

Figure 51 Change Ownership of a Recipient Dataset or Target Group



10.5.1 Changing Ownership of Recipient Datasets and Lookup Tables

To change ownership of a specific set of Recipient Datasets and Lookup Tables, click on **Recipients Datasets** from the main Recipient Warehouse screen. A screen will be displayed that shows a list of all groups and accounts that have at least one Recipient Dataset or Lookup Table in their Recipient Warehouse. Click on an individual account or group name to see the datasets and lookup tables owned by that account or group.

Figure 52 The Recipient Datasets and Lookup Tables Screen

Recipient Datasets and Lookup Tables

This page displays an overview of all groups and accounts which have at least one recipient dataset or lookup table in their recipient warehouse. Click on an owner name to open the administration page for datasets and lookup tables of that owner.

Owner Name	Type	Recipient Datasets	Lookup Tables
brandnew	Group	1	2
demo	Group	5	10
landover	Group	11	10
marketing	Group	2	4
newdemo	Group	2	3
sales	Group	4	7
senate-test	Group	8	8
tc	Group	2	4

From the Recipient Dataset and Lookup Table Administration screen, change ownership of Recipient Datasets and Lookup Tables from one owner (either an individual or group) to another. To assign new ownership, select a new owner from the drop-down menu, and click **[OK]**. This will assign all of the Recipient Datasets and Lookups tables to the new owner. To delete all Datasets and Lookup Tables for the selected account or group, click **[Delete]**.

Figure 53 The Recipient Dataset and Lookup Table Administration Screen

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 *demo*

Select a New Owner:

Datasets Owned: 4

Lookup Tables Owned: 9

Below you see a list of all datasets and lookup tables owned. Click on a dataset or unreferenced lookup table to select it.

Recipient Datasets	Lookup Tables	References
L-Soft international, Inc.	Countries	<not referenced>
PetsMart Company	Courses	<not referenced>
Widget Company	Industry	1 field of 1 recipient dataset or mailing list
XYZ Company	Models	<not referenced>
	Pets	1 field of 1 recipient dataset or mailing list
	States	1 field of 1 recipient dataset or mailing list
	XYZ Products	1 field of 1 recipient dataset or mailing list
	departments	1 field of 1 recipient dataset or mailing list
	transportation	<not referenced>



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 (i.e. all of the datasets and lookup tables belonging to that owner).

From the Recipient Dataset and Lookup Table Administration screen, select individual datasets or unreferenced lookup tables for deletion. Clicking on one of the Recipient Datasets opens the Recipient Dataset Administration screen. From that screen, view details for the individual dataset. The Recipient Dataset Administration screen lists dataset name, a description of the

dataset, the current owner, and the number of Hosted Lists in the dataset. From this screen, it is possible to **[Delete]** the dataset or **[Cancel]** to return to the previous screen.

Figure 54 The Recipient Dataset Administration Screen

Recipient Dataset Administration

Click the Delete Button to remove the selected dataset.

Recipient Dataset: **XYZ Company**
 Description: XYZ Customer Center
 Current Owner: Group **demo**
 Hosted Lists in Dataset: 4

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

Hosted Lists
DEVELOP-XYZ
NEWS-FROM-XYZ
TALK-BACK-TO-XYZ
XYZ Monthly Newsletter

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 55 The Hosted List Administration Screen

Hosted List Administration

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

Hosted List: **XYZ Monthly Newsletter**
 Description: monthly customer newsletter
 Current Owner: Group **demo**
 Recipient Dataset: **XYZ Company**

From the Recipient Dataset and Lookup Table Administration screen, it is possible to click on individual tables that are not currently referenced. Clicking on individual table names opens the Lookup Table Administration screen. From this screen, view the name of the Lookup Table, a description of the table, and the current owner of the table. From here, it is possible to **[Delete]** the Lookup Table or **[Cancel]** to return to the previous screen.

Figure 56 The Lookup Table Administration Screen

Lookup Table Administration

Click the Delete Button to remove the selected lookup table.

Lookup Table: **State_Abbr**
 Description: This lookup table contains the two letter abbreviation of each state in the United States
 Current Owner: Group **sales**

10.5.2 Changing Ownership of Recipient Target Groups

To change ownership of Recipient Target Groups, click on **Recipient Target Groups** from the Recipient Warehouse screen. This will open the Change Target Group Owner screen. From here, select an individual account or group in order to view all Recipient Target Group categories owned by that group.

Figure 57 The Change Target Group Owner Screen



Click on an individual category to view details for that category. From here, view the category name and current owner. Assign a new owner by selecting an account or group from the drop-down menu and clicking **[OK]**. It is also possible to **[Delete]** the current category or **[Cancel]** to return to the previous screen. On the lower half of the screen, a list of Target Groups contained within the currently selected category, along with descriptions of those Target Groups appears. Click on an individual Target Group to change ownership of a particular Target Group, or to delete the Group.

Figure 58 Changing Ownership of a Target Group Category



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 **Global Component Setting > Administration Hub > 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 LISTSERV Maestro Test-Bed Backup](#).

Figure 59 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:

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 the **Test** link, 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 popup window before the process terminates, will not stop the process; it will continue running. 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.



Important: Commands must not define external processes that run indefinitely. Each external process should terminate itself when it has completed the action it is supposed to execute. Processes that run continuously slow the server down over time. Eventually

this 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 **Administration Hub**, and then **General Administration**. The General Component Settings for Administration Hub screen opens.
- For the Maestro User Interface component, click **Maestro User Interface**, and then **General Administration**. The General Administration of Maestro User Interface screen opens.
- For the Maestro Tracker component, click **Maestro Tracker**. The General Component Setting for Maestro Tracker screen opens.



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 3 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 such as a backup tape or other 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 that was saved in step 3.

- **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 that was saved in step 3.



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. 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.

7. 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.

Section 12 Using a LISTSERV Maestro 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 **Global Component Settings > Administration Hub > 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](#) 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 a 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.

- 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 do not configure the `Port` entry, then the default port (80 for HTTP or 443 for HTTPS) will be used.

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:

- If the HTTP port is changed on a server that houses the Maestro User Interface component, then it is also necessary to change the External Host Name setting for Hosted Datasets and Hosted Lists in the Administration Hub. Do so by editing the Hosted Data Settings on either the global component level or the group or single user level. See Section 7 [Settings for the Maestro User Interface](#).
- If the HTTP port is changed on a server that houses the Maestro User Interface component, then it is also necessary to make some changes to the INI file of the Maestro User Interface component. On the server with the Maestro User Interface installed, locate the `lui.ini` file:

```
[maestro_install_folder]/lui/lui.ini
```

Edit or add the entry `ExternalHTTPPort`. (If the entry is not defined, it defaults to port 80 if HTTP is used and port 443 if HTTPS is used.) See Section 20 [Editing LISTSERV Maestro INI Files](#).

- If the HTTP port is changed on the server that houses the Maestro Tracker component, then it is also necessary to change the Tracking URL – HTTP Port settings in the Administration Hub by editing the **Global Component Settings** and any group or single user level settings that are defined. See Section 7 [Settings for the Maestro User Interface](#).

- If the HTTP port is changed on the server that houses the Administration Hub component, then it is necessary to edit the Maestro User Interface INI file. On the server where the Maestro User Interface component is installed, locate the `lui.ini` file: `[maestro_install_folder]/lui/lui.ini`

Edit or add the entry “HubExternalHTTPPort”. (If the entry is not defined, it defaults to port 80 if HTTP is used and port 443 if HTTPS is used.) See Section 20 [Editing LISTSERV Maestro INI Files](#).

- 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/lui
```

- 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/lui.host
```

- The “`hub.host`” file is located on the server where the Administration Hub component is installed, while the `lui.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/lui
```

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/lui
```

- 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.
- 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.



Important: Maestro Tracker will work most effectively if it uses port 80. Many sites have firewalls that prevent their users from connecting to other ports for HTTP connections, which would not only prevent them from being tracked, but from reaching the actual Web pages whose access is being tracked.

14.2.2 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`.

14.2.3 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 the **Global Component Settings** icon, and then **Maestro Tracker**. Edit the port number. Click **[OK]** to save. The change will be effective immediately.

14.2.4 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.

14.2.5 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 IP Addresses

By default, LISTSERV Maestro binds the HTTP port on all IP addresses of the server on which it is running. 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 HTTP port by using any of the server's addresses.

No changes to the LISTSERV Maestro configuration are required if this default behavior is satisfactory. However, to make LISTSERV Maestro bind to only a single IP address on the server, you need to edit the `BindAddress` entry in the Tomcat INI file:

```
[maestro_install_folder]/conf/tomcat.ini
```

Example: `BindAddress=192.168.1.1`



Notes: If several LISTSERV Maestro components are installed on the same server, then all of them will be affected by this change. 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.

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.

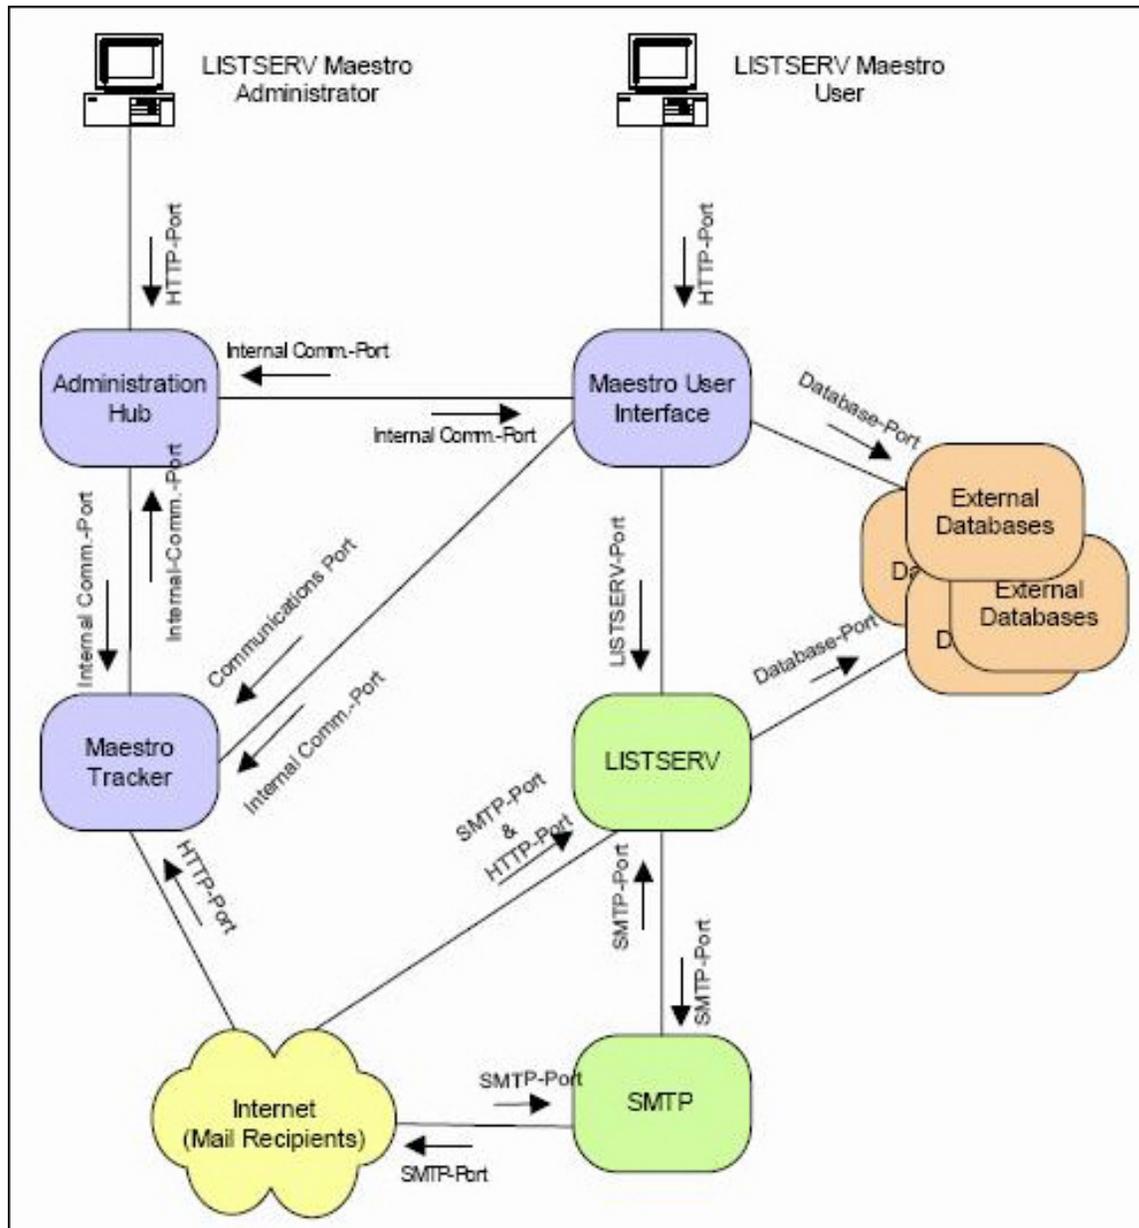
Communication occurs exclusively using TCP ports (see the Section 14 [Using Non-Standard Ports](#) for more information). If the components are installed behind, in front of, or on both sides of a firewall, then the firewall needs to be configured to let communication through on certain ports between certain servers. Figure 60 shows LISTSERV Maestro components and all other players (the Maestro Administrator, the Maestro User, and the Internet, which represents the messages recipients) and their interconnections.

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**. (This does not apply for the HTTP Port between the Internet and Maestro Tracker, which can never be replaced by the HTTPS Port).
- **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 60 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 Components** icon. Click **Maestro User Interface**, and then **General Administration**. The General Administration of Maestro User Interface screen opens. Check **Disallow multiple logins with the same user account**.

Figure 61 Multiple Logins

General Administration of Maestro User Interface

General Settings

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

Event transfer interval (from Maestro Tracker): min.

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

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

Subscriber access page translations are automatically retrieved by the system during startup. Click the button below to refresh the page translations now, during runtime.



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
```

Figure 62 XML Entry for the Session Timeout

```
<!-- 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
```

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.

LISTSERV Maestro also offers the possibility of using SSL for communication with the Administration Hub and/or the Maestro User Interface components. 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?

Only the Maestro User Interface and/or the Administration Hub component should be secured with SSL - never the Maestro Tracker component. The Maestro Tracker component always requires use of normal HTTP; it cannot be configured to use HTTPS (because the collection of the tracking events needs to be fast, HTTPS is too slow for this).

As all components installed on one server share the same access method, it is necessary to select the access method for all components simultaneously. Therefore, if it is desired to secure the Administration Hub and/or Maestro User Interface components with SSL, they must be installed on a separate server (or separate servers) from the Maestro Tracker component (however, both may be on the same server, as long as they are both secured). Similarly, if only one component is to be secured (either the Administration Hub or the Maestro User Interface), then each component has to be secured on a separate server. This will permit security of one independent of the other.

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 web site:

<http://java.sun.com/j2se/1.4/docs/tooldocs/win32/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 63 Example of Base64 Encoded Outfile

```
-----BEGIN NEW CERTIFICATE REQUEST-----
tPnJhsLOuocsBYAmyM1lqiZ5BEVWAnJfZ6kyN/Xft5NFxGIy9Uynz5kODfBwFUgiu98iQKWyMK
C/bGFuZ2VuMQ8wDQYDVQKKEwZMLVnVZnQxEDAObgNVBAsTB1Vua25vd24xDzANBgNVBAMTBnRl
cHBp6E7Zyl9wkPyVpnlqbnbtXQGAablJIInE9/LruaJ1NX1f/NVJgL4vPiDKsU4laGvJHBNhdj+
F0uVb3SIb3DQEBAUAA4GBAB6XqdfJvhy7dTHijsHjw+c4ELQFI/TkHBvgp5NaCccQoNwwW9ln
IeOikDb2lwG6LiYfpVBss5+OOW2jXlq9CdNw1KLSdQ+kMtZjdVr8+iQ9gsqxvskCAwEAAa
AAMA0GCSqGMIIBpjCCAQ8CAQAwZjELMAkGA1UEBhMCREUxEDAObgNVBAsTB0dlcm1hbnkxETAP
BgNVBAcTCEVYyZCBnzANBkgqhkiG9w0BAQEFAAOBjQAwYkCgYEAz+hQRsqDWRLvmV4YD5+JaQ
EXn5qqJeyzkfg2PQoU2VPgHID0VnyTpt8r/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 the 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 HOME\j2sdk1.4.0_01\jre\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 61.

Figure 64 Imported Certificate

```
Owner: OU=For VeriSign authorized testing only. No assurances (C)VS1997,
OU=www.verisign.com/repository/TestCPS Incorp. By Ref. Liab. LTD.,
O="VeriSign, Inc"
Issuer: OU=For VeriSign authorized testing only. No assurances (C)VS1997,
OU=www.verisign.com/repository/TestCPS Incorp. 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 any other value, 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).

Finally, you may have to edit the `Port` entry in the same `tomcat.ini` file. This entry specifies the port to be used for receiving client requests.

If this entry is not specified, the appropriate default port will be used, i.e. in case of HTTPS (which you have enabled by specifying `"SecureServer=true"`) the default HTTPS-port 443 would be used.

However, if the `Port` 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 `Port` entry is present, does it specify the correct port you want to use even for HTTPS? 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 `Port` entry, please see also Section 14.2.1 [Configuring the HTTP Port](#) about other issue that may apply when changing the HTTP port (or in this case the HTTPS port).

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: It is normally not a good idea to secure access to Maestro Tracker; therefore, if you want to secure your Administration Hub and/or Maestro User Interface components with SSL, then you **must** install them on a separate server from the Maestro Tracker component (however, those two may be on the same server, as long as you plan to secure both of them).

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 4 [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, as shown in Figure 65.

Figure 65 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, as shown in Figure 66.

Figure 66 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 affected component in order to take effect.

Table 4 Maestro User Interface INI-File Entries

Entry Key	Description
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 23.1 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 23.2 Allowing or Disallowing Bi-Directional Character Sets .
ChangeLog	Specifies whether or not a change log of subscriber activities (for hosted lists and datasets) is to be created. (Optionally) If a change log file is used, then this defines the time period for each log (i.e. daily, weekly, monthly, or yearly). See Section 13.2 Subscriber Activity Change Log . Default: <code>false</code>
ClickThroughURL	Path-part of the click-through tracking URL used for URLs without passing of merged parameters. Default: <code>/trk/click</code>

Entry Key	Description
ClickThroughPPURL	Path-part of the click-through tracking URL used for URLs with passing of merged parameters. Default: /trk/clickp
CreatePersonalProfileTables	Specify whether full personal profiles should be stored in the system database when personal tracking is used. Default: true See Section 19 Tracking and Recipient Profiles
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 23.1 Defining the Default Mail Charset
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 cleans them out). Default: 1 See the LISTSERV Maestro Data Administrator's Manual for more information about subscription confirmations.
ExternalHostName	The host name of the Maestro User Interface server as seen by clients that access the user interface with a web browser. To be used if clients shall access the server with a different host name than the normal internal host name of the server, for example when using a proxy. Default: The value of the <code>HostName</code> parameter, if present. Otherwise the normal host name (canonical host name) of the server running the Maestro User Interface. See Section 21.3.1 See Section 21 Distributed Components .

Entry Key	Description
ExternalHTTPPort	<p>The HTTP (or HTTPS) port of the Maestro User Interface server as seen by clients that access the user interface with a web browser. To be used if the external port of the server is different than the default for the protocol employed (both if the normal HTTP or the secure HTTPS protocol is employed).</p> <p>Default: 80 (if the protocol is HTTP) or 443 (if the protocol is HTTPS)</p> <p>See Section 14.2.1 Configuring the HTTP Port</p>
ExternalProtocolIsHTTPS	<p>Defines if HTTPS is used as the access protocol for the Maestro User Interface web-interface pages. Set to “true” if HTTPS is used, or “false”, if normal HTTP is used.</p> <p>Default: false</p> <p>See Section 18 Securing Access With SSL</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>
HubExternalHostName	<p>The host name of the Administration Hub server as seen by clients that access the hub interface with a web browser. To be used if clients shall access the hub server with a different host name than the normal internal host name of the server, for example if the server has several names or when using a proxy.</p> <p>Default: the host name defined by the RegistryHubHost parameter. If that parameter is not present either, the normal host name (canonical host name) of the server running the Maestro User Interface.</p> <p>See Section 21 Distributed Components</p>
HubExternalHTTPPort	<p>The HTTP (or HTTPS) port of the Administration Hub server as seen by clients that access the hub user interface with a web browser. To be used if the external</p>

Entry Key	Description
	<p>port of the server is different than the default for the protocol employed (both if the normal HTTP or the secure HTTPS protocol is employed).</p> <p>Default: 80 (if the protocol is HTTP) or 443 (if the protocol is HTTPS)</p> <p>See Section 14.2.1 Configuring the HTTP Port</p>
HubExternalProtocolIsHTTPS	<p>Defines if HTTPS is used as the access protocol for the Administration Hub web-interface pages. Set to “true” if HTTPS is used, or “false”, if normal HTTP is used.</p> <p>Default: <code>false</code></p> <p>See Section 18 Securing Access With SSL</p>
HubRMIPort	<p>Internal communication port (RMI-Port) of the Administration Hub server.</p> <p>Default: 1099</p> <p>See Section 14.2.2 Configuring the Internal Communication Port</p>
MaintenanceMode	<p>Defines if the Maestro User Interface component will run in maintenance mode or not.</p> <p>Default: <code>false</code></p> <p>See Section 8.3 Runtime Administration and System Shutdown</p>
OpenUpURL	<p>Path-part of the open-up tracking URL.</p> <p>Default: <code>/trk/open</code></p>
OptInConfirmExpiration	<p>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.</p> <p>Default: 48</p> <p>See the LISTSERV Maestro Data Administrator’s Manual for more information about subscription confirmations.</p>
RegistryHubHost	<p>Host name of the server with the Administration Hub component.</p> <p>Default: <code>localhost</code></p>
RemoteAdminPassword	<p>Password for remote log file access.</p> <p>Default: <code>none</code> (no remote log file access allowed)</p>

Entry Key	Description
RestoreBackup	<p>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.</p> <p>Default: <i>none</i></p> <p>See Section 11.7 Restoring a Backup</p>
RMIPort	<p>Internal communication port (RMI-Port) of the Maestro User Interface server.</p> <p>Default: 1099</p> <p>See Section 14.2.2 Configuring the Internal Communication Port</p>
TrackerHost	<p>Host name of the server with the Maestro Tracker component.</p> <p>Default: localhost</p> <p>See Section 21.2.3 Moving the Maestro Tracker Component to Another Server</p>
TrackerRMIPort	<p>Internal communications port (RMI-Port) of the Maestro Tracker server</p> <p>Default: 1099</p> <p>See Section 14.2.2 Configuring the Internal Communications Port</p>
UnsubscribedCleanupInterval	<p>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.)</p> <p>Default: 10</p> <p>See the LISTSERV Maestro Data Administrator’s Manual for more information about subscription confirmations.</p>

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 affected component to take effect.

Table 5 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.
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 14.2.2 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 affected component to take effect.

Table 6 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

Entry Key	Description
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 14.2.2 Configuring the Internal Communication Port
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 14.2.2 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 affected 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 7 Basic Configuration for Tomcat INI-File Entries

Entry Key	Description
BindAddress	The local IP-address that the server shall bind to (i.e. the server will accept connections only if directed to this IP Address). Default: By default, the server binds to <u>all</u> addresses of the local computer (i.e. accepts connections on all of these addresses). See Section 15 Defining IP Addresses .

Entry Key	Description
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 (or HTTPS) port used to accept client connections. Default: 80 (for HTTP) or 443 (for HTTPS).
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)	
SecureServer	Defines if the server is supposed to be secure (<code>true</code>) or not (<code>false</code>). 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". 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 filemane must include the full absolute path to the file (for Windows, include the drive letter). Default: <code>[install_folder]/conf/keys.keystore</code>
KeystorePassword	Only used if "SecureServer=true". The password used to access the keystore-file 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 8 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
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 25 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

Entry Key	Description
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>
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>

Section 21 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 “distribute 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.

21.1 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).

21.2 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.

21.2.1 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 below). Remember to skip steps 2, 6, and 8 below, which are only applicable when using the internal database.
 - 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 4.1 [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 **Global Component Settings > Administration Hub > General Administration**, and then click 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 2 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.

21.2.2 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.

Also, if the new Administration Hub server has an external name and/or port which differs from the externally known name (i.e. the name you used for “RegistryHubHost”), then you also might need to add the entries “HubExternalHostName” and/or “HubExternalHTTPPort”. See Section 20.4 [Tomcat INI-File Entries](#) for details.

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.

21.2.3 Moving the Maestro Tracker Component to Another Server

Important: 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 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 averse 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.
6. Restart LISTSERV Maestro (on all servers), as well as any LISTSERV instances that have been stopped.
7. If the new server of the Maestro Tracker component is now known with a different name than the previous server was, you also might need to adjust the “Tracking Host” setting in the Administration Hub. For this, log into the Administration Hub and set the correct tracking host either on global default level, group level or user level, whichever is applicable.

21.2.4 Moving the Database External Component to Another Server

1. Install the database software on the new server and start it.
2. Follow the instructions in Section 4.1 [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.).
3. 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.

21.3 Server Name Aliases and Proxies

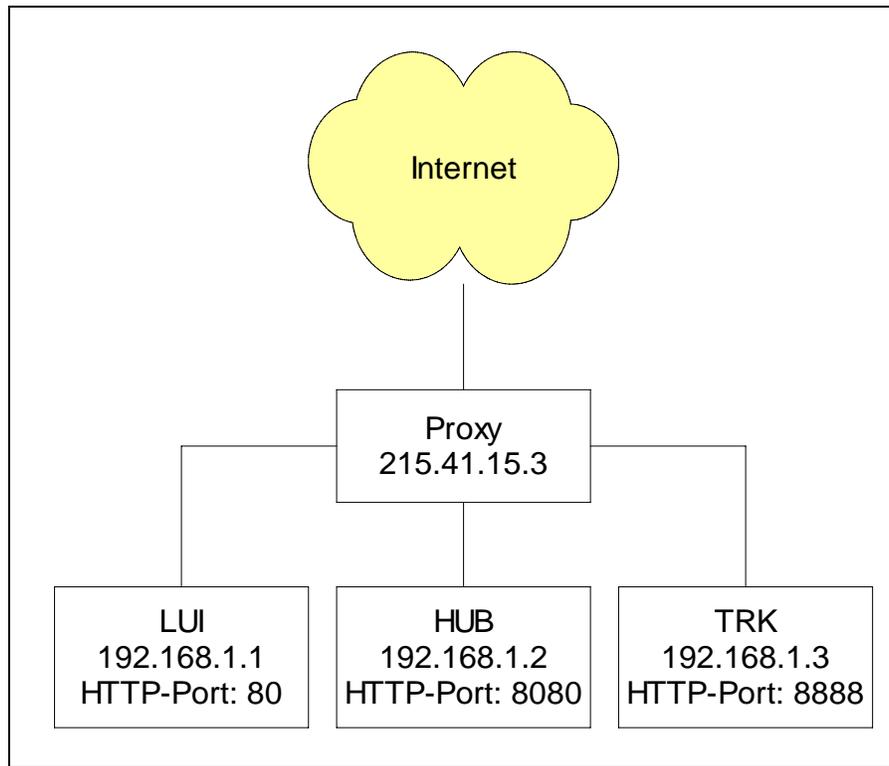
With any given installation of LISTSERV Maestro, the components of LISTSERV Maestro are installed on one or more servers, where each server has its own host name. Components on separate servers use the other server or servers' name(s) to access the component(s) there. Similarly, the "outside" world (users 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).

Figure 677 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 67 Sample Proxy Setup

**Example 1**

Assume that the proxy shown in Figure 677 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/lui`.

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.

21.3.1 Configuring LISTSERV Maestro Components with Server Name Aliases or Proxies

If the local name or HTTP port of a host differs from the externally known name or port the following files must be edited:

- If the **Maestro User Interface** component’s server has an external host name that is different from the local name, edit the file:

```
[maestro_install_folder]/lui/lui.ini
```

Add or edit the entry “ExternalHostName” to read:

```
ExternalHostName=HOST
```

where “HOST” is replaced with the external name of the server running the Maestro User Interface component. Do not include port information in the “ExternalHostName” parameter.

- If the **Maestro User Interface** component's HTTP (or HTTPS) server has an external port number that is different from the default port number (port 80 for HTTP, port 443 for HTTPS), edit the file:

```
[maestro_install_folder]/lui/lui.ini
```

Add or edit the entry "ExternalHTTPPort" to read:

```
ExternalHTTPPort=PORT
```

where "PORT" is replaced with the external HTTP (or HTTPS) port number of the server running the Maestro User Interface component.

- If the **Administration Hub** component's server has an external host name that is different from the internal name, edit the file:

```
[maestro_install_folder]/lui/lui.ini
```

Add or edit the entry "HubExternalHostName" to read:

```
HubExternalHostName=HOST
```

where "HOST" is replaced with the external host name of the server running the Administration Hub component. Do not include port information in the "HubExternalHostName" parameter.

- If the **Administration Hub** component's HTTP (or HTTPS) server has an external port number that is different from the default (port 80 for HTTP, port 443 for HTTPS), edit the file:

```
[maestro_install_folder]/lui/lui.ini
```

Add or edit the entry "HubExternalHTTPPort" to read:

```
HubExternalHTTPPort=PORT
```

where "PORT" is replaced with the external HTTP (or HTTPS) port number of the server running the Administration Hub component.



Important: The entries above must go into the `lui.ini`, not into the `hub.ini`.

- If the **Maestro Tracker** component's server has an external host name or HTTP port that is different from the internal name/port, set the external host name and/or port in the **Tracking URL** settings in the Administration Hub. Please see Section 7.3 [Setting the Default Tracking URL](#) for more details.

To carry out the examples above, the following changes to the administration settings would have to be made:

- For **Example 1**, four `lui.ini` entries are required:

```
ExternalHostName=maestro.sample.com
```

```
ExternalHTTPPort=9001
```

```
HubExternalHostName=maestro.sample.com
```

```
HubExternalHTTPPort=9002
```

In addition, the Administration Hub would be used to configure the tracking URL to use a Tracker Host of `maestro.sample.com` and a HTTP Port of 9003.

- For **Example 2**, two `lui.ini` entries are required:

```
ExternalHostName=lui.sample.com
```

```
HubExternalHostName=hub.sample.com
```

In addition, the Administration Hub would be used to configure the tracking URL to use a Tracker Host of `trk.sample.com` and a HTTP Port of 80. Next, it would be necessary to configure the proxy accordingly, so that it transparently forwards the requests as described above – but this depends on the proxy software used and is not part of the LISTSERV Maestro setup.

Section 22 LISTSERV Maestro in Evaluation Mode

In order to function in normal mode, LISTSERV Maestro needs to be connected to a fully licensed instance of LISTSERV. LISTSERV Maestro 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 version. If not, then LISTSERV Maestro 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, including job definition, authorization, and viewing delivered jobs, without actually being able to use LISTSERV Maestro 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 in evaluation mode.

In addition, for test delivery in evaluation mode to function, the following condition must be fulfilled: Either LISTSERV’s SMTP listener or an instance of LSMTP that is connected to LISTSERV must be installed on the same server that the LISTSERV instance used for delivery is running on. In addition, this SMTP listener/LSMTP must be listening for SMTP requests on the standard SMTP port 25.

Section 23 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.

23.1 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 9 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)

Charset Name:	Description:
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 <u>not</u> 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`.

23.2 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 24 Authenticating Message Origin with DomainKeys Signatures

LISTSERV Maestro allows you to 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 allows you to 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 Component Settings** icon, click **Maestro User Interface**, and then click **Default DomainKeys Settings**. The DomainKeys Settings screen opens, allowing you to define the default behavior for DomainKeys signatures.

Figure 68 DomainKeys Settings Screen

DomainKeys Settings

Define if e-mails that are sent by the application shall be signed for DomainKeys by default. Additionally, decide on application default level if the users are allowed to decide about DomainKeys signatures (using the default value defined above).

Note: These settings are only applicable if the LISTSERV host of the account or group supports DomainKeys.

Default Setting

No, don't sign e-mails for DomainKeys

Yes, sign e-mails for DomainKeys

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

OK Cancel

This screen also lets you define whether or not users are allowed to change the DomainKeys signature settings for each job.

To setup the DomainKeys signature settings for a group and user, click on the **Administer User Accounts** icon, click on the name of a group or user, then click **DomainKeys Settings**. In addition to the options described above, this page also allows you to define whether or not a certain group or single user is supposed to use the application-wide defaults or special local 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, sign e-mails for DKIM** and **The user must use the setting supplied above without changes for each mail job** on the DomainKeys Setting screen (for application-wide defaults), and then make sure that all groups/users have the **Use inherited value** option selected.

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 25 Adding 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 `script`, 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 (for example 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 25.2 [Defining the Default Context](#) for details).

25.1 Adding Content as a New Context

To create a new context, create a folder (with the name of the context) in the `webapps` folder of LISTSERV Maestro. For example, if you want to create a new context called `sample`, then you would create the following folder:

```
[install_folder]\webapps\sample\
```

Copy the following folder (and the files in it) from the `archives` context to your newly created context:

```
[install_folder]\webapps\archives\WEB-INF\
```

(The `archives` context is automatically installed with each Maestro Tomcat.

After you copy this folder into your new folder, you will have:

```
[install_folder]\webapps\sample\WEB-INF\
```

Open the `WEB-INF` folder, and then open the `web.xml` file with a text editor (i.e. Notepad). The file contains a `<filter>` definition with two parameter definitions. The first parameter is the context parameter and is defined as follows:

```
<param-name>context</param-name>
<param-value>archives</param-value>
```

Change the parameter’s value (which is currently “archives”) to the name of your own context. For example, if you called your context “sample”, then:

```
<param-name>context</param-name>
<param-value>sample</param-value>
```

You can now put whatever files you want into the `sample` folder (you can also create subfolders). You may want to include a start page, such as `index.html` or similar, but you can also have other pages, image files, or other downloadable content if you want (even in subfolders). These files would then be accessed as follows:

```
http://yourhost/sample           access to the start page (if one was supplied)
http://yourhost/sample/         also access to the start page
http://yourhost/sample/page.html access to page.html
http://yourhost/sample/sub/other.html access to other.html in the sub
                                   subfolder
```

Restart Maestro to make it aware of the new context.



Important: 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 `script`).

Do not put any files into the `WEB-INF` folder, as they would not be accessible via a URL.

Do not create a folder or subfolder called `META-INF`. This is a reserved name.

25.2 Defining the Default Context

The default context is the context used if the user types the access URL without any context name, such as `http://yourhost` or `http://yourhost/somepage.html`. This is not supported out of the box by Maestro; instead, it must be configured.

To define a default context, create a new context (with a new subfolder in “webapps”) as described in the [previous section](#).

In addition, before restarting Maestro, include the following entry in the `tomcat.ini`:

```
DefaultContext=CONTEXT_NAME_HERE
```

where you replace `CONTEXT_NAME_HERE` with the name of your context (for example, `DefaultContext=sample`, making the `sample` context the default).



Note: Do not define any of the reserved LISTSERV Maestro contexts as the default context (`lui`, `hub`, `trk`, `list`, `archives`, or `script`)

With this entry in the `tomcat.ini`, the given default context (“`sample`” in this case) can now be accessed in two ways:

- As before, you can still specify the context directly in the URL:

```
http://yourhost/sample           access to the start page
http://yourhost/sample/         also access to the start page
http://yourhost/sample/page.html access to page.html
http://yourhost/sample/sub/other.html access to other.html 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://yourhost</code>	access to the start page of the default context
<code>http://yourhost/</code>	also access to the start page of the default context
<code>http://yourhost/page.html</code>	access to <code>page.html</code> in the root folder of the default context
<code>http://yourhost/sub/other.html</code>	access to <code>other.html</code> in the sub subfolder of the default context

Section 26 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.

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.

Glossary of Terms

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.

Bounce – An Email message that is returned as undeliverable

Bounce Server – An optional dedicated LISTSERV server used exclusively to process bounced email.

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.

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.

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.

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”.

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.

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.

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.

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”.

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.

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.

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.).

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.”

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.

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.

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.

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.

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.

Appendix A: LISTSERV Maestro Standard Default Ports

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

Index

A

Accessing	
HUB	4
log files	92
restricting access to components	105
Admin Password	
changing	11
Administration	
configuring backups	51
email notifications	54
introduction	1
policies	51
refreshing the Subscriber Page translations	56
runtime administration	51
special administrative user account	68
system shutdown	51
Administrative Policies	51
configuring backups	51
email notifications	54
user restrictions	52
Archived Jobs	
importing	70
Archiving	
delivered jobs	68

B

Backups	
configuring	51
configuring a backup time	80
configuring external post-backup processes	81
configuring the backup history	83
configuring the backup location	82
creating a Test-Bed Backup on the original system	90
ID 84	
restoring	80, 84
restoring a Test-Bed Backup into the Test System	90
saving	80
saving to an external medium	84
using a Test-Bed Backup	88

C

Character Sets	
allowing bi-directional charsets	151
defining the default mail charset	149
disallowing bi-directional charsets	151
using	149
Configuring	
backups	51
LISTSERV Maestro for first use	9
port usage	97
the application server shutdown port	100
the backup history	83
the backup location	82

the backup time	80
the external database	13
the external post-backup processes	81
the HTTP port	97
the internal communication port	99
the internal database connection port	100
the tracker communications port	100

D

Database Plugins	21
IBM DB2 v7.2	21
IBM DB2 v8.2	21
MySQL ConnectorJ Driver	22
MySQL L-Soft Driver	22
ODBC Driver	25
Oracle 8i, 9i, 10g Thin Driver	23
registering	26
SQL Server i-net SPRINTA Driver	24
SQL Server jTDS Driver	23
SQL Server Microsoft Driver	24
Datasets	
changing ownership	76
Distributed Components	
configuring LISTSERV Maestro components with	
server name aliases or proxies	145
fresh installation	137
introduction	136
moving components to another server	137
moving the Administration Hub component	139
moving the database external component	143
moving the Maestro User Interface component	137
moving the Tracker component	141
server name aliases and proxies	143
DomainKeys	
using	152
Drop-In Content Elements	
changing ownership	74

E

Email Notifications	
administration	54
testing	55
Evaluation Mode	148
External Database	<i>See</i> System Database
adding	19
configuring	13
defining connections	20
removing	19

F

Firewall	
install behind	102

H

HTTP Ports	
configuring	97
HUB Interface	
accessing	4
changing a password	11
editing the INI file	130
introduction	4
ports	96

I

Importing	
archived jobs	70
INI Files	
Administration HUB entries	130
editing	125
Maestro User Interface entries	126
Tomcat entries	132
Tracker entries	131
Internal Database	<i>See System Database</i>
adding	19
removing	19
International Character Sets	<i>See Character Sets</i>
using	149
IP Addresses	
defining	101

J

Jobs	
archiving	68
changing ownership	70
importing archived jobs	70

L

LISTSERV	
preparing for LISTSERV Maestro	28
preparing Maestro to send distribute jobs	36
preparing to allow HLL	30
preparing to process distribute jobs from Maestro	28
specifying a separate LISTSERV instance for processing bounces	38
specifying the LISTSERV host with different internal and external names	37
using existing lists with Maestro	38
with LISTSERV Maestro	28
LISTSERV Maestro	
configuring for first use	9
default settings	43
editing INI files	125
in evaluation mode	148
installing behind a firewall	102
introduction	1
LISTSERV instances	28

log files	92
making aware of the server certificate	120
multiple tracking URLs	50
preparing LISTSERV for	28
preparing LISTSERV to allow HLL	30
preparing LISTSERV to process distribute jobs	28
preparing to send distribute jobs to LISTSERV	36
securing access with SSL	111
setting the default tracking URL	49
settings for the User Interface	40
specifying a separate LISTSERV instance for processing bounces	38
standard default ports	162
User Roles	2
using existing LISTSERV lists	38
LISTSERV Web Interface	
adding to the Tomcat server	156
Log Files	
accessing	92
Remote Log Access	92
Subscriber Activity Change Log	93
Lookup Tables	
changing ownership	76

M

Maestro Tracker	
configuring the Tracker communications port	100
editing the INI file	131
ports	96

P

Password	
changing	11
Ports	
configuring the application server shutdown port	100
configuring the HTTP Port	97
configuring the internal communication port	99
configuring the internal database connection port	100
configuring the tracker communications port	100
configuring usage	97
for the HUB	96
for the TRK	96
for the User Interface	96
standard defaults	162
using non-standard ports	96

R

Recipient Datasets	
changing ownership	76
Recipient Profiles	122
Remote Version Query	
HUB	3
LUI	3
TRK	3

Reports	
changing ownership	70
Restrictions	
disallowing concurrent access with the same user	
account	107
IP Address.....	105
to components.....	105
Runtime Administration	
settings	51

S

Secure Communication	
components	114
creating an unsigned server certificate	115
installing a server certificate	115
installing a signed server certification.....	118
installing a trusted root certificate.....	118
introduction.....	111
making Maestro aware of the server certificate	120
performing a CSR	117
securing access with SSL.....	111
securing the Trusted Root Certificate Keystore	115
Sender Profiles	
changing ownership	74
SSL	
securing access.....	111
Subscriber Pages	
refreshing the translations	56
System Database	
adding an external DB	19
available plugins	21
configuring the external DB.....	13
DB2.....	17
defining external database connections	20
introduction.....	12
MySQL	17
optimization	18
Oracle.....	16
preparing	15
removing an internal DB.....	19
SQL Server	16
System Requirements	
Client	3
System Shutdown	51

T

Target Groups	
changing ownership	76, 79
Test-Bed Backups	88
creating	90
restoring into the Test System.....	90

Tomcat	
adding content	154
adding content as a new context	154
adding the LISTSERV Web Interface	156
defining the default context	155
editing the INI file	132
Tracker	
configuring the communications port.....	100
editing the INI file	131
ports.....	96
Tracking	122

U

User Accounts	
assigning single user settings.....	58
creating	57
deleting	67
disallowing concurrent access	107
editing.....	57, 58
special administrative account.....	68
User Interface	
editing the INI file	126
ports.....	96
setting the default tracking URL.....	49
settings.....	40
User Restrictions	52
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.....	43

W

What's New	
Automatic Bounce Processing.....	xii
Creating a Test-Bed Backup.....	xiv
Customization Wizard.....	x
DomainKeys Signatures	xiv
Flexible Dataset Administration	xii
introduction	x
LUI Admin Account.....	xiii
Refreshing the Subscriber Page Translations	xiii
Resource Translation Tool	xi
Shared Subscriber Management and Reporting.....	xii
Subscriber Options	xi
System Support	xiv
Tracking Links in Drop-Ins	xiii
Tracking Restrictions	xiii

This page intentionally left blank.
