



Whitepaper

Understanding How LISTSERV® Handles Bounces

August 25, 2010

Copyright © 2010 L-Soft international, Inc.

Information in this document is subject to change without notice. Companies, names, and data used for example herein are fictitious unless otherwise noted. Some screen captures have been cropped and/or edited for emphasis or descriptive purposes.

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

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

Copyright © 2010, L-Soft international, Inc.

All Rights Reserved Worldwide.

LISTSERV is a registered trademark licensed to L-Soft Sweden and L-Soft international, Inc.

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

Introduction

Everyone who runs an electronic mailing list faces the problem of handling invalid e-mail addresses. These "bad addresses" generate a "bounce," often referred to as a Delivery Status Notification (DSN) or Non-delivery notification (NDN), for every message sent to the list. High-traffic mailing lists can result in a very large number of bounces, necessitating action on the part of the list owner to prevent server slow-down and other negative consequences.

Handling bounces generally is the responsibility of the person or people running the mailing list. LISTSERV has some valuable tools that make it easier to manage bounces and take some of the burden off of the list owner. However, due to the nature of e-mail and network communication, none of these tools can be 100 percent effective.

The purpose of this white paper is to provide you with a better understanding of how LISTSERV handles bounces and some of the issues involved. In the interest of keeping this document concise, some technical issues have been simplified. For more in depth information and a wider perspective, see the resources referred to in the final section of this document.

E-Mail Protocols and Bounces

E-mail protocols are a set of standards that determine how e-mail gets from one place to another. Two of the most important protocols that govern this behavior are *RFC 821 Simple Mail Transfer Protocol* (SMTP for short) and *RFC 822 Standard For The Format Of ARPA Internet Text Messages*. Often, these protocols are referred to only by their RFC number, which is short for Request for Comments. The full text of RFCs can be found at <http://www.ietf.org/rfc.html>. References are included at the end of this document.

RFC 821 spells out the protocol for how e-mail gets from one place to another. Contained within this protocol is the return path address for messages that bounce. This address is called the MAIL FROM address and it is extremely important because it instructs the receiving mail server where to send back any non-delivery notices that are generated as a result of the transaction.

RFC 822 spells out the format of text messages that are sent over the Internet. Contained within this protocol is the format for message headers. Part of the information contained in message headers is to whom the message is being sent (TO:) and whom the message is from (FROM:).

Throughout this document a distinction is made between the address used in the MAIL FROM: field and the one used in the FROM: field of the same message. The key point is that the MAIL FROM: address does not have to be the address of the person who sent the message, and does not have to be the FROM: address that appears in the headers. Often times the MAIL FROM: address and the FROM: address are two completely different addresses, as is the case when messages are distributed by a LISTSERV-hosted mailing list.

How LISTSERV Uses the MAIL FROM: and From: Addresses

When a message is distributed by LISTSERV, the From: address will be the address of the person who originally sent the message to the list. The MAIL FROM: address will be an address handled by LISTSERV, something similar to `owner-listname@peach.ease.lsoft.com`. The reason for using these addresses is so that subscribers to the list can see who sent the message, but any bounces that are generated by the mailing go to the owner-address rather than the original sender of the message (who would not be in a position to do anything about them).

This idea can be clarified by drawing a comparison to postal mail ("snail mail"). When someone sends you a letter, it comes in an envelope that includes the address of the person who sent the letter and your address. When you receive the letter you open and discard the envelope then read the message itself. The message may also contain information about whom the message is from and who it is to, and this information may differ from the information that was on the envelope. If the letter cannot be delivered for some reason, it is marked "Return To Sender" (or similar) and returned to the sender's address on the envelope.

Internet mail standards are based on a similar idea. The MAIL FROM: address is equivalent to the return address on the envelope. The actual message headers and the body of the message are equivalent to the letter itself, and they should not be referred to at all when the message is being delivered (though sometimes broken mail servers do make reference to them).

LISTSERV's Automatic Bounce Detection

Messages can bounce for many reasons. Messages that bounce contain error codes that define the characteristics of why the message bounced. (The codes are documented in RFC 821, RFC 1893 "*Enhanced Mail System Status Codes*" and RFC 1894 "*An Extensible Message Format for Delivery Status Notifications*") These reasons are separated into two categories, "soft bounces" (or temporary) and "hard bounces" (or permanent). Soft bounces reflect a temporary condition for why the message could not be delivered, for example the recipient's mailbox was full. Hard bounces reflect a permanent condition for why the message could not be delivered such as the address does not exist.

As stated earlier, LISTSERV always sends out list messages with a MAIL FROM: pointing to `owner-listname@[your site]`. Any bounces generated by a mailing go to the list's bounce-processing address. LISTSERV monitors the owner-listname address and assumes any mail that is delivered there is a bounce.

Automatically Deleting List Addresses That Return "Hard" Bounces

LISTSERV has the ability to automatically delete addresses on the list that return hard bounce error codes, saving the list owner time and effort. By using the `Auto-delete` keyword and setting it to `Full-Auto` or `Semi-Auto`, you are telling LISTSERV to attempt to decipher the bounced message by reading the error code. LISTSERV looks for "hard bounces" (permanent errors like no such user, no such host, and so on). If the failing recipients are subscribed to the

list, LISTSERV removes them and notifies the list owner. No action is required from the list owner.

This works reasonably well, but there are a few pitfalls:

- Bounces are not always sent to the correct address. According to Internet standards, any bounces generated are supposed to go to the MAIL FROM: address, but there are operating mail servers that do not fully comply with the standards. Some bounces can go to virtually any e-mail address in the headers of the message: the From:, the Sender:, the Reply-To:, and so on.

Occasionally, bounces are even sent to the mailing list itself. Fortunately, LISTSERV has ways to detect when this is happening so that the bounces are not distributed to the mailing list. Instead, suspected bounces are forwarded to the owner.

A more pernicious problem is that some bounces will go to addresses that have nothing to do with LISTSERV. LISTSERV cannot process messages it never sees, so these types of bounces will have to be dealt with manually. Since this is the result of mail servers you have no control over, there is not much that can be done about this problem apart from recognizing that it exists.

- LISTSERV recognizes bounces that are in standard formats. Specifically, it recognizes bounces in the formats described in RFC 821 and in RFC 1894. LISTSERV also recognizes bounces in "LMAIL format" (a proprietary format developed for L-Soft products before RFC 1894 was available). Unfortunately, there are many non-standard formats being used that LISTSERV does not recognize. Some non-delivery notifications do not even say what e-mail address resulted in the bounce. Any message LISTSERV does not understand is forwarded to the list owner. This is true even if you have `Auto-delete=Yes, Full-Auto` set (unless you use probing, see the following section).
- Someone may subscribe to your mailing list with one address and then have that address forward mail to a second address. For instance, someone may sign up as `joeuser@bigfoot.com` and then have that address forward mail to `j.user@aol.com`. If something goes wrong with the second address, LISTSERV will be notified that the second address is bad, but it will not be able to do anything, even if the bounce is in a recognizable format, because the second address is not subscribed (the first address is).

Depending on your `Auto-delete=` settings, LISTSERV will either forward such bounces to the list owner or will just add the address to the daily monitoring report.

`Auto-delete=Yes, Semi-Auto` does the former, `Auto-delete=Yes, Full-Auto` the latter. This situation can also be addressed by probing.

- Problems may arise if the bounce error codes are used incorrectly, especially confusing temporary (soft) and permanent (hard) bounce codes. For instance, some sites report "Mailbox full"-type errors as 550 errors, (a code for a permanent error) and will cause the address to be deleted from the list even though the problem will be cleared up as soon as the user cleans out his or her mail account. (Most sites will use the correct code for this, which is 552.)

The key point is that as far as the standards are concerned, and so as far as LISTSERV is concerned, the numerical code is the only thing that matters, regardless of what any qualifying text may say. Sites using the codes incorrectly may see their users incorrectly getting removed from mailing lists. All the list owner can do is add the affected subscribers back and inform them of what the problem is that caused them to be deleted from the mailing list.

Passive Probing

One of the banes of large, active mailing lists is the large volume of bounces that LISTSERV receives but cannot interpret and therefore must be handled manually by the list owner. A tool called "passive probing" allows LISTSERV to handle most bounces, even if they are not in a standard format. This is how it works:

Depending on the list's configuration, a certain number of subscribers is randomly chosen during each mailing to receive the message, but with a customized MAIL FROM: address. This type of message is called a "probe." The purpose of the probe is to find out if the address is valid. If the address is not valid, (it was returned to the special MAIL FROM: address) LISTSERV can take appropriate action, regardless of whether the actual bounce is in an understandable format or came from a forwarded address.

This is a more concrete example:

A mailing list has an address of "listname@home.ease.lsoft.com". One of the subscribers has an address of badaddress@example.com. If this subscriber was selected to receive a probe, LISTSERV would set the MAIL FROM: address to:
owner-listname*badaddress**example*-com@home.ease.lsoft.com

If the subscribed badaddress@example.com is bad, a bounce is generated and is sent to the MAIL FROM: address (in this example, to owner-listname*badaddress**example*-com@home.ease.lsoft.com). LISTSERV sees that a message was sent to that address and extracts the name of the mailing list ("listname" in this case) and the offending address, which is badaddress@example.com. LISTSERV then acts according to the Auto-delete= settings.

This is a very powerful tool and using it results in a very substantial decrease in the amount of time the list owner has to spend dealing with bounces. (Active probing uses the same mechanism to detect bounces on renewals.) There are some disadvantages to be aware of:

- Mail servers that do not send to the correct address will continue to be a problem.
- LISTSERV can now process bounces that are in non-standard formats, but all messages sent to the probe address that LISTSERV cannot interpret will be treated as hard bounces. This includes auto-responses, Mailbox Full-type messages, and non-delivery notifications warning of a temporary delivery problem.
- The MAIL FROM: addresses that are used to probe are completely legal in that they fully comply with the relevant standards. However, they are not typical addresses, and some mail servers may reject mail with such an unusual MAIL FROM: address. If this appears to be happening, it is usually because the mail server has been configured to reject mail from addresses containing asterisks ('*'s), though it could also be because of the length of the address. Often if you explain what probing is and convince the e-mail administrator that you are not doing anything wrong, they will modify the rules to fix this, but of course there is no guarantee.

Where to Go for More Information

Both the List Owner's Manual and the Site Manager's Manual for LISTSERV contain a lot of information on auto deletion and active and passive probing that is not discussed here. The manuals are available at: <http://www.lsoft.com/resources/manuals.asp>

Internet standards relating to SMTP and those relating to the format of e-mail messages, including headers, can be found at <http://www.ietf.org/rfc.html>

For the purposes of this document, the most important standards are:

- RFC 821, "*Simple Mail Transfer Protocol*" <http://www.ietf.org/rfc/rfc0821.txt>
- RFC 822, "Standard For The Format Of ARPA Internet Text Messages" <http://www.ietf.org/rfc/rfc0822.txt>
- RFC 1123, "*Requirements for Internet Hosts -- Application and Support*" (in particular chapter 5, "Electronic Mail -- SMTP and RFC-822") <http://www.ietf.org/rfc/rfc1123.txt>
- RFC 1893, "*Enhanced Mail System Status Codes*" <http://www.ietf.org/rfc/rfc1893.txt>
- RFC 1894, "An Extensible Message Format for Delivery Status Notifications" <http://www.ietf.org/rfc/rfc1894.txt>