University of Tennessee
5 June 2002

Recommendations for Automatic Responses to Electronic Mail

draft-moore-auto-email-response-00

Status of this Memo

This document is not currently associated with any working group. Comments on this internet-draft should be sent to the mailing list <ietf-822@imc.org>, or to the author. Such comments should cite the Internet-Draft identifier draft-moore-auto-email-response-00 so others can be sure you are commenting on the same version they read.

Abstract

This memo makes recommendations for software that automatically responds to incoming electronic mail messages, including "out of the office" response generators, mail filtering software, email-based information services, and other automatic responders. The purpose of these recommendations is to discourage undesirable behavior which is caused or aggravated by such software, to encourage uniform behavior (where appropriate) among automatic mail responders, and to clear up some sources of confusion among implementors of automatic email responders.

1. Introduction

Many programs which automatically respond to email are currently in use. Although these programs vary widely in their function, several problems with this class of programs have been observed, including: significant numbers of useless or unwanted response and responses sent to inappropriate addresses, and occasional incidences of mail loops or "sorcerer's apprentice" syndrome. This memo recommends behavior for programs that automatically respond to electronic mail in order to reduce the number of problems caused by such programs.

1.1 Types of automatic responses

There are several different types of automatic responses. At least two types of automatic responses have been defined in IETF standards - Delivery Status Notifications [1] which are intended to report the status of a message delivery by the message transport system, and Message Disposition Notifications [2] which are intended to report of the disposition of a message after it reaches a recipient's mailbox. These responses are defined elsewhere and are generally not within the purview of this document, except that this document recommends specific cases where they should or should not be used.

Other types of automatic response in common use include:

- "Out of office" or "vacation" notices, which are intended to inform the sender of a message that the message is unlikely to be read, or acted on, for some amount of time;
- Email-based information services, which accept requests (presumably from humans) via email, provide some service, and issue responses via email also. (Mailing lists which accept subscription requests via email fall into this category);
- Information services similar to those mentioned above except that they are intended to accept messages from other programs;
- Various kinds of mail filters (including "virus scanners") which act on behalf of a recipient to alter the content of messages before forwarding them to that recipient, and issue responses in the event a message is altered; and
- Responders designed to filter unsolicited messages from programs (e.g. a program that responds to any message from an unknown or unverifiable source and requires that party to "demonstrate signs of intelligent life" before the original message can be read.)

Recognizing the wide variety of response types in use, these recommendations distinguish between several classes of automatic responders according to the party or service on whose behalf the responder acts:

- "Service Responders" exist to provide access to some service via email requests and responses. These are permanently associated with an email address, and when sending to such an address the sender presumably expects an automatic response. An email-based file retrieval service is an example of a Service Responder.
- "Personal Responders" exist to make automatic responses on behalf of a single human recipient, in advance of, or in lieu of, that recipient reading the message. These responders operate according to criteria specified by the individual recipient. The UNIX "vacation" program is an example of a Personal Responder.
- "Group Responders" exist to make automatic responses on behalf of any of a group of human recipients, in advance of, or in lieu of, a response from the actual recipient. Group Responders are similar to Personal Responders except that in the case of a Group Responder the criteria for responding are not set by the individual

recipient. A "virus scanner" program that filtered all mail sent to a group of recipients (say, every recipient in a particular DNS domain) and sent responses when a message was rejected or delivered in an altered form, would be an example of a Group Responder.

Appropriate behavior for a responder varies from one class to another. A behavior which might be appropriate from a Service Responder (where the sender is expecting an automatic response) might not be appropriate from a Personal Responder. For example, a Service Responder might send a very long response to a request, or one that is not in a human-readable format, according to the needs of that service. However a Personal Responder should assume that a human being is reading the response and send only brief responses in plain text.

1.2. Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [3].

2. Format of automatic responses

The following sections specify details of the contents of automatic responses, including the header of the response message, the content of the response, and the envelope in which the response is transmitted to the email transport system.

2.1 Message header

The fields in the message header SHOULD be set as follows:

2.1.1 From field

In correspondence between humans, the From field serves multiple purposes: It identifies the author of the message (or in some cases, the party or parties on whose behalf the message was sent), and it is the default destination of replies from humans. Also, unfortunately some mail systems still send nondelivery reports and other kinds of automatic responses to the From address.

For automatic responses, the role of the From field in determining the destination of replies from humans is less significant, because in most cases it is not useful or appropriate for a human (or anyone) to reply to an automatic response. (The exception is when there is some problem with the response; it should be possible to provide feedback to the person operating the responder).

So the From address in an automatic response needs to be chosen according to the following criteria:

- To provide an indication of the party or agent on whose behalf the response was sent,
- To provide an address to which a recipient of an inappropriate response can request that the situation be corrected, and

- To diminish the potential for mail loops.

The following behavior is thus recommended:

- For responses sent by Service Responders, the From field SHOULD contain an address which can be used to reach the (human) maintainer of that service, and the human-readable portion of the From field (the phrase preceding the address) SHOULD contain a name or description of the service to identify the service to humans.
- For responses sent by Personal Responders, the From field SHOULD contain the name of the recipient and an address chosen by the recipient to be recognizable to correspondents. Normally this would be the same address that was used to send mail to that recipient.

  In the case of a recipient having multiple mail addresses forwarded to the same mailbox (and responder), a Personal Responder MAY use heuristics to guess, based on the information available in various message header fields, which of several addresses for that recipient the sender is likely to have used, and use that address in the From field of the response. However any address chosen by this method MUST have been explicitly allowed by the recipient on whose behalf the responder is operating.

  Note: Due to privacy reasons it may be inappropriate for responders to disclose an address that is derived, say, from the recipient's login information (e.g. POP or IMAP user name or account name on a multiuser computer) or which discloses the specific name of the computer where the response was generated. Furthermore these do not necessarily produce a valid public email address for the recipient. For this reason the From field of a Personal Response SHOULD be settable by the recipient on whose behalf the responder is acting.
- For Group Responders, the From address SHOULD contain an email address which could be used to reach the maintainer of that Group Responder. Use of the Postmaster address for this purpose is NOT RECOMMENDED.

  The human-readable portion of the From address (the phrase before the address) SHOULD contain an indication of the function performed by the Group Responder and on whose behalf it operates (e.g. "Example Agency virus filter")

2.1.2 To field

The To header field SHOULD indicate the recipient of the response. In general there SHOULD only be one recipient of any automatic response. This minimizes the potential for sorcerer's apprentice syndrome and denial-of-service attacks.

2.1.3 Date field

The Date header field SHOULD indicate the date and time at which the response was composed. This MUST NOT be taken as any indication of the delivery date of the subject message, nor of the time at which the response was sent.

2.1.4 Subject field

The Subject field SHOULD contain a brief indication that the message is an automatic response, followed by contents of the Subject field (or a portion thereof) of the subject message. The prefix "Auto-Re:" MAY be used as such an indication.

NOTE: Just as the prefix "Re:" (presumably an abbreviation of the English word "reply") is sometimes translated to other languages by mail readers, or otherwise interpreted by mail readers as indication that the message is a reply, so the prefix "Auto-Re:" may also be translated or used as a generic indication that the message is an automatic response. However the "Auto-Re:" indication is intended only as an aid to humans in processing the message. The validity of "Auto-Re:" SHOULD NOT be assumed by mail processing software.

2.1.5 In-Reply-To field

The In-Reply-To field SHOULD be included in the header of the response message if there was a Message-ID field in the subject message. If present in the response, the In-Reply-To field SHOULD contain the message-id of the subject message. A References field MAY also be supplied.

2.1.6 Auto-Submitted field

The Auto-Submitted field, with a value of "auto-replied", SHOULD be included in the message header of any automatic response. See section 5.

2.2 Message content

In general, messages sent by Personal or Group Responders SHOULD be brief, and in text/plain format. A multipart/alternative construct MAY be used to communicate responses in multiple languages if it is desirable to use multiple charsets.

Response messages SHOULD NOT include significant content from the subject message. In particular responses SHOULD NOT contain non-text/plain attachments from the subject message.

2.2.1 Use of DSNs and MDNs for automatic responses

An exception to the above policy can be made for responders whose purpose is to filter out harmful content from incoming email. In such cases it may be appropriate to issue a delivery status notification (DSN) or a message disposition notification (MDN) to indicate that such mail has been refused, deleted, or altered. Such a responder MAY issue a DSN if the responder is operating as a part of the mail transport system and has access to the message envelope, and the response is generated on or prior to delivery to the recipient's mailbox. Alternatively, a response MAY use the MDN format, provided the response is generated on or after delivery to a recipient's mailbox. An MDN SHOULD NOT be issued as an automatic response unless the subject message contains a Disposition-Notification-To field. In all cases such responses MUST conform to the DSN or MDN specifications.

For example, in the case of a DSN, the Action per-recipient field SHOULD be set to "failed" with a Status code of 5.7.1 (Delivery not authorized, message refused) if the message was not delivered due to security reasons, and the Action field SHOULD be set to "relayed" or "delivered" (as appropriate) with a Status code of 2.6.4 (conversion with loss performed) if the message was modified to remove significant (presumably harmful) content before relay or delivery but the remainder of the message was relayed or delivered to its destination.

In the case of an MDN, a disposition mode of "automatic-action/-MDN-sent-automatically" would be appropriate, with a disposition-type of "deleted" or "denied" with a disposition modifier of "error" for messages which were automatically discarded, and a disposition-type of "processed" with a disposition modifier of "warning" for messages which were filtered before being presented to the recipient. The Failure: or Warning: MDN fields could be used to supply additional information about the reason for refusal or alteration of the message.

2.3 Message envelope

The SMTP MAIL FROM address, or other envelope return address used to send the message, SHOULD be chosen in such a way as to make mail loops unlikely. A loop might occur, for instance, if both sender and recipient of a message each have automatic responders - the recipient's responder sends mail to the sender's responder, which sends mail back to the recipient's responder.

The primary purpose of the MAIL FROM address is to serve as the destination for delivery status messages and other automatic responses. Since in most cases it is not appropriate to respond to an automatic response, and the responder is not interested in delivery status messages, a MAIL FROM address of <> MAY be used for this purpose. A MAIL FROM address which is specifically chosen for the purpose of sending automatic responses, and which will not automatically respond to any message sent to it, MAY be used instead of <>.

The RCPT TO address should be the address of the intended recipient of the response. It is RECOMMENDED that the NOTIFY=NEVER parameter of the RCPT command be specified if the SMTP server supports the DSN option [4].

3. When to send automatic responses

An automatic responder MUST NOT send a response for every message received. In practice there are always reasons to refuse to respond to requests. The criteria for deciding whether to respond will differ from one responder to another, according to the responder's purpose. In general, care should be taken to avoid sending useless or redundant responses, and to avoid contributing to mail loops and facilitating denial-of-service attacks.

Here are some broad guidelines:

- Automatic responses SHOULD NOT be issued in response to any message which contains an Auto-Submitted header field with a value of "auto-replied" or "auto-generated".

- Personal and Group responses whose purpose is to notify the sender of a message of a temporary absence of the recipient (e.g. "vacation" and "out of the office" notices) SHOULD NOT issue the same response to the same sender more than once within a period of several days, even though that sender may have sent multiple messages. A 7-day period is RECOMMENDED as a default.
- Personal and Group responses whose purpose is to notify the sender of a message of a temporary absence of the recipient (e.g. "vacation" and "out of the office" notices) SHOULD NOT be issued unless a valid address for the recipient is explicitly included in the To, CC, or Bcc field of the subject message. Since a recipient may have multiple addresses forwarded to the same mailbox, recipients SHOULD be able to specify a set of addresses to the responder which it will recognize as valid for that recipient.
- Responders SHOULD NOT generate responses for any null address. Responders MAY refuse to generate responses for addresses commonly used as return addresses by responders - e.g. those with local-parts matching "owner-*", "*-request", "MAILER-DAEMON", etc. Responders SHOULD check the destination address for validity before generating the response, to avoid cluttering up the local mail queues with messages that cannot be delivered or are unlikely to be useful.
- In order to avoid responding to spam and to certain kinds of attacks, automatic responses from Service Responders should be sent only for well-formed requests. This may include checking that the message resulting in the response has a content-type and content appropriate to that service.

4. Where to send automatic responses (and where not to send them)

In general, automatic responses SHOULD be sent to the address given in the Return-Path field, or if the responder has access to the message envelope, the reverse-path from the SMTP MAIL command, or (in a non-SMTP system) another envelope return address which serves as the destination for nondelivery reports.

If the Return-Path field is not present in the subject message, there is a bug in the SMTP server that delivered the message, or that SMTP server is improperly configured. A Personal or Group responder SHOULD NOT deliver a response to any address other than that in the Return-Path field, even if the Return-Path field is missing. It is better to fix the problem with the mail delivery system than to rely on heuristics to guess the appropriate destination of the response.

A Service Responder MAY deliver the response to the address from the From field, or to another address from the request payload, provided this behavior is precisely defined in the specification for that service. The Reply-To field SHOULD NOT be used for this purpose.

The Reply-To field SHOULD NOT be used as the destination for automatic responses from Personal or Group Responders. In general, this field is set by a human sender based on his/her anticipation of how human recipients will respond to the specific content of that message. Even for replies from humans, there are cases where it is not appropriate to respond to the Reply-To address, especially if the sender has asked that replies be sent to a group and/or mailing list. Since a Personal or Group Responder operates on behalf of a human recipient, it is safer to assume that any Reply-To field present in the message was set by a human sender on the assumption that any reply would come from a human who had some understanding of the roles of the sender and other recipients. An automatic responder lack the information necessary to understand those roles. Sending automatic responses to Reply-To addresses can thus result in a large number of people receiving a useless or unwanted message; it can also contribute to mail loops.

Use of the From field as the destination for automatic responses has some of the same problems as use of Reply-To. In particular, the From field may list multiple addresses, while automatic responses should only be sent to a single address. In general, the From and Reply-To addresses are used in a variety of ways according to differing circumstances, and for this reason Personal or Group Responders cannot reliably assume that an address in the From or Reply-To field is an appropriate destination for the response.

Similarly, the Sender field SHOULD NOT be used as the destination for automatic responses. This field is intended only to identify the person or entity that sent the message, and is not required to contain an address that is valid for replies.

The Return-Path address is really the only one from the message header that can be expected, as a matter of protocol, to be suitable for automatic responses that were not anticipated by the sender.

5. The Auto-Submitted header field

The purpose of the Auto-Submitted header field is to indicate that the message was originated by an automatic process, or an automatic responder, rather than by a human; and to facilitate automatic filtering of messages from signal paths for which automatically generated messages and automatic responses are not desirable.

5.1 Syntax

The syntax of Auto-Submitted is as follows:

```
auto-submitted-field    = "Auto-Submitted:" CFWS
                          auto-submitted [CFWS] CRLF

auto-submitted          = ( "no" / "auto-generated" /
                          "auto-replied" / extension )
                              opt-parameter-list

extension               = token

opt-parameter-list      = *( [CFWS] ";" [CFWS] LWSP parameter )
```

The symbols "token", and "parameter" are as defined in [5].

5.2 Semantics

The Auto-Submitted header field SHOULD NOT be supplied for messages that were manually submitted by a human. Such a field MAY be supplied for a manually sent message that is intended to test the response of other mail system components to the presence of an Auto-Submitted field in a message.

The auto-generated keyword:

- SHOULD be used on messages generated by automatic (often periodic) processes (such as UNIX "cron jobs"),
- MUST NOT be used on manually generated messages,
- MUST NOT be used on a message issued in direct response to another message.

The auto-replied keyword:

- SHOULD be used on messages sent in direct response to another message,
- MUST NOT be used on manually-generated messages,
- MUST NOT be used on messages generated by automatic or periodic processes.

The "no" keyword may be used to explicitly indicate that a message was originated by a human.

Extension keywords may be defined in the future, though it seems unlikely. The syntax and semantics of such keywords must be published as RFCs and approved using the IETF Consensus process [6]. Keywords beginning with "x-" are reserved for experiments and use among consenting parties.

Optional parameters may also be defined by an IETF Consensus process. The syntax of optional parameters is given here to allow for future definition should they be needed. Implementations of Auto-Submitted conforming to this specification MUST NOT fail to recognize an Auto-Submitted field and keyword that contains syntactically valid optional parameters, but such implementations MAY ignore those parameters if they are present. Parameter names beginning with "x-" are reserved for experiments and use among consenting parties.

The "comment" syntactical construct can be used to indicate a reason why this message was auto-submitted.

6. Security Considerations

Automatic responders introduce the possibility for several kinds of attack, including:

- Use of such responders to relay harmful or abusive content (worms, viruses, spam, and spymail) for the purpose of wider distribution of the content or masking the source of such content;
- Use of such responders to mount denial-of-service attacks by using responders to relay messages to large numbers of addresses, or to flood individual mailboxes with a large amount of unwanted content, or both;
- Deliberate or accidental use of such responders to construct mail loops or "sorcerer's apprentice syndrome", thus taxing the resources of the mail transport system;

- In addition, the responder itself may be subject to attack by sending it large numbers of requests.

This document attempts to reduce the vulnerability of responders to such attack, in particular by

- Recommending that responders not relay significant content from the subject message (thus minimizing the potential for abusive content)
- Recommending that responders clearly mark responses with the "Auto-Submitted: auto-replied" header field to distinguish them from messages originated by humans (in part, to minimize the potential for loops and denial-of-service attacks),
- Recommending that Personal and Group Responders limit the number of responses sent to any individual per period of time (also limiting the potential damage caused by loops),
- Recommending that responders respond to at most one address per incoming message (to minimize the potential for deliberate or accidental denial-of-service via "multiplication" or sorcerer's apprentice syndrome),
- Recommending that responses should be brief and in plain text format (to minimize the potential for mail responders to be used as mechanisms for transmitting harmful content and/or disguising the source of harmful content).

However, because email addresses are easily forged, attacks are still possible for any email responder which does not limit access and require authentication before issuing a response. The above measures attempt to limit the damage which can be done, but they cannot entirely prevent attacks.

This section describes vulnerabilities inherent in automatically responding to mail. Other vulnerabilities are associated with some mail-based services which automatically respond to email messages, but these are not caused by the fact that the server automatically responds to incoming messages. In general, all network based services (including those accessed by email) need to provide security that is sufficient to protect the resources that are accessible by the service against inappropriate use.

7. IANA Considerations

Section 5 of this document defines two new extension mechanisms - new keywords for the auto-submitted header field, and new optional parameters for the auto-submitted field. If at any point in the future new keywords or paramters are approved (through an IETF Consensus process) it may be appropriate for IANA to create a registry of such keywords or paramters.

8. Acknowledgments

In the mid-1990s Jeroen Houttuin of TERENA authored a series of internet-drafts on "Behavior of Mail Based Servers", and in particular, one document on "Answering Servers" [7]. While these documents were (to this author's knowledge) never formally published, they provided the first well-reasoned argument (known to this author) as to the best way for such servers to interface with email systems and protocols.

The idea for the auto-submitted field comes from the X.400/MHS mail system [8]. [9] defined an "Autosubmitted" field for use when gatewaying between X.400 and Internet mail. Jacob Palme wrote an internet-draft [10] defining use of the "Auto-Submitted" field for Internet mail, which made it through Last Call without significant objections, but got stalled in an attempt to resolve non-substantial objections. The definition of Auto-Submitted in this document is derived (i.e. slightly simplified) from the one in that document, with some text stolen outright.

Thanks are also due to those who contributed suggestions to this document: (so far) Eric Hall, Florian Weimer, and Dan Wing.

9. Author's Address

Keith Moore
Innovative Computing Laboratory
University of Tennessee, Knoxville
1122 Volunteer Blvd, #203
Knoxville, TN 37996-3450

moore@cs.utk.edu

10. References

[1] Moore, K. Vaudreuil, G. An Extensible Message Format for Delivery Status Notifications. RFC 1894, January 1996. (non-normative reference)

[2] Fajman, R. An Extensible Message Format for Message Disposition Notifications. RFC 2298, March 1998. (non-normative reference)

[3] Bradner, S. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119, March 1997. (normative reference)

[4] Moore, K. SMTP Service Extension for Delivery Status Notifications. RFC 1891, January 1996. (normative reference, but only barely)

[5] Freed, N. Borenstein, N. Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies. RFC 2045, November 1996. (normative reference)

[6] Narten, T., Alvestrand, H. Guidelines for Writing an IANA Considerations Section in RFCs. RFC 2434, October 1998. (normative reference)

[7] Houttuin, J. BoMBS series: Behavior of Mail Based Servers / Part 2: A-BoMBS / Answering Servers. Expired Internet-Draft draft-rare-msg-a-bombs-01.txt, December 1994. Available at http://google.com/ (non-normative reference, work apparently no longer in progress, included only for attribution)

[8] X.400. (perhaps someone can supply the correct reference for the first version of the X.400 document to define autosubmitted?) (non-normative reference)

[9] Kille, S. MIXER (Mime Internet X.400 Enhanced Relay): Mapping between X.400 and RFC 822/MIME. RFC 2156, January 1998. (non-normative reference)

[10]
Palme, J. "The Auto-Submitted and Expires Headers in E-mail". Expired Internet-Draft "draft-ietf-mailext-new-fields-15.txt", February 1999. (non normative reference, work apparently no longer in progress, included only for attribution)