Internet-Draft
Expires: 2 March 2004

K. Moore
University of Tennessee
2 September 2003

Recommendations for Automatic Responses to Electronic Mail

draft-moore-auto-email-response-03

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
`http://www.ietf.org/1id-abstracts.html`

The list of Internet-Draft Shadow Directories can be accessed at
`http://www.ietf.org/shadow.html`

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-03 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" or "vacation" 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.

Intended status: Once it appears that this document has received sufficient review, comment, and community support, the author intends to submitted it as an individual submission for Proposed Standard status. Proposed Standard seems more appropriate than BCP because this document describes protocols more than operational practices.

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" mode. This memo recommends behavior for programs that automatically respond to electronic mail in order to reduce the number of problems caused by such programs.

(Note: the term "sorcerer's apprentice mode" is defined as a bug in a protocol where, under some circumstances, the receipt of a message causes multiple messages to be sent, each of which, when received, triggers the same bug.) (From [I1])

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 [I2] which are intended to report the status of a message delivery by the message transport system, and Message Disposition Notifications [I3] 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,
- "Change of address" notices, intended to inform the sender of a message that the recipient address he used is obsolete and that a different address should be used instead (whether or not the subject message was forwarded to the current address),
- "Challenges", which require the sender of a message to demonstrate some measure of intelligence and/or willingness to agree to some conditions before the subject message will be delivered to the recipient (often to minimize the effect of "spam" or viruses on the recipient),
- 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, and
- 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.

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 one or more email addresses, 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. A calendar service that allowed appointment requests to be made via email, and which responded to such requests, would be another example of a Service Responder.
- "Personal Responders" exist to make automatic responses on behalf of a single recipient address, in advance of, or in lieu of, that recipient reading the message. These responders operate according to criteria specified on a per-recipient basis. The UNIX "vacation" program is an example of a Personal Responder. A responder that accepts mail sent to a single address, attempts to analyze and classify the contents, and then issues a response which is dependent on that classification, is also a Personal Responder.
- "Group Responders" exist to make automatic responses on behalf of any of a significant set of recipient addresses (say, every recipient in a particular DNS domain), 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 on a per-recipient basis. A "virus scanner" program that filtered all mail sent to any recipient on a particular server, and sent responses when a message was rejected or delivered in an altered form, might 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 and Definitions

The key words "MUST", "MUST NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", and "MAY" in this document are to be interpreted as described in [N1].

The term "subject message" is used to refer to a message which causes a response to be sent.

The term "response" refers to a message that is automatically issued on receipt of a subject message by a responder.

A "responder" is a process that automatically responds to subject messages under some well-defined set of conditions.

Unless specified otherwise, the term "recipient" refers to the email addresses to which a subject message was delivered (rather than, for instance, the address to which the response was sent). A "recipient" address might be permanently associated with a

responder, or it might be the address of a human being whose mail is, under some conditions, answered by a responder.

2. When (not) 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 some kinds of received messages, e.g. for loop prevention, to avoid responding to "spam", to avoid being used as a means to launder or amplify abusive messages, to avoid inappropriately revealing personal information about the recipient (e.g. to avoid an automatic indication that a recipient has not read his mail recently), and to thwart denial-of-service attacks against the responder. 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 or 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 (see below), where that field has any value other than "no".
- Personal and Group responses that are intended to notify the sender of a message of the recipient's inability to read or reply to the message (e.g. "away from my mail" or "too busy" notifications) 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 a recipient (e.g. 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.

  Note: RFC 2822 section 3.6.3 permits varying uses of the Bcc field, some of which would allow the sender of the subject message to explicitly specify the recipient's address as a "Bcc" recipient without a Bcc field appearing in the message as delivered, or without the Bcc field in the delivered message containing the recipient's address. However, perhaps because Bcc's are rarely used, the heuristic of not responding to messages for which the recipient was not explicitly listed in a To, CC, or Bcc header field has been found to work well in practice.
- Personal and Group Responders MAY refuse to generate responses except to known correspondents or addresses of otherwise "trusted" individuals. Such responders MAY also generate different kinds of responses for "trusted" vs. "untrusted" addresses. This might be useful, for instance, to avoid inappropriate disclosure of personal information to arbitrary addresses.

- Responders SHOULD NOT generate any response for which the destination of that response would be a null address (e.g. an address for which SMTP MAIL FROM or Return-Path is <>), since the response would not be delivered to a useful destination. 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 are encouraged to check the destination address for validity before generating the response, to avoid generating responses 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 NOT be sent for extremely malformed requests. This may include checking that the subject message has a content-type and content appropriate to that service.
- Because the vast majority of email is unauthenticated, and return addresses are easily forged, in order to avoid being used as a means of denial-of-service attacks (i.e. to flood mailboxes with unwanted content) Service Responders SHOULD NOT return large responses (say, more than a few kilobytes) without specific knowledge that the request was actually authorized by the party associated with the address to which the response will be sent. Similarly, Service Responders SHOULD NOT cause unwanted side-effects (such as subscribing the sender to a mailing list) without reasonable assurance that the request was authorized by the affected party.

  NOTE: Since each responder has a different purpose and a different set of potential threats to which it might be subjected, whether any particular means of authentication is appropriate for a particular responder is not in scope for this document.
- A responder MAY refuse to send a response to a subject message which contains any header or content which makes it appear to the responder that a response would not be appropriate. For instance, if the subject message contained a Precedence header field [I4] with a value of "list" the responder might guess that the traffic had arrived from a mailing list, and would not respond if the response were only intended for personal messages. For similar reasons, a responder MAY ignore any subject message with a List-* field [I5]. (Because Precedence is not a standard header field, and its use and interpretation vary widely in the wild, no particular responder behavior in the presence of Precedence is recommended by this specification.)

3. 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.

3.1 Message header

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

3.1.1 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. 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 to the response 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. One exception is when there is some problem with the response; it should be possible to provide feedback to the person operating the responder.

So in most cases 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. The human-readable portion of the From field (the display-name 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. Often this will be the same address that was used to send the subject message 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 it MUST be possible for a recipient on whose behalf the responder is acting to explicitly specify the human-readable name and address to be used in the From header fields of responses.

  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 MUST 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, see [N2], section 3.2.6) SHOULD contain an indication of the function performed by the Group Responder and on whose behalf it operates (e.g. "Example Agency virus filter")

3.1.2 Reply-To field

If a reply is expected by the responder, the Reply-To field of the response SHOULD be set to the address at which the reply is expected, even if this is the address of the same or another responder. Responders which request replies to be sent to responders MUST prevent mail loops and sorcerer's apprentice mode. Note that since (according to the previous section) the From field of the response SHOULD contain the address of a human, if the Reply-To field of the response is used to direct replies to a responder it will not be the same as the address in the From field.

Discussion: this assumes that the human recipient's user agent will normally send replies to the Reply-To address (if present), as recommended by [I6] since 1982, but that it is still possible for a recipient to reply to the From address if he or she finds it useful to do so. This is consistent with the intended use of these fields in [I6] and [N2].

3.1.3 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 mode and denial-of-service attacks.

3.1.4 Date field

The Date header field SHOULD indicate the date and time at which the response was generated. 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.

3.1.5 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) from the subject message. The prefix "Auto:" MAY be used as such an indication. If used, this prefix SHOULD be followed by an ASCII SPACE character (0x20).

NOTE: Just as the (Latin-derived) prefix "Re:" that is commonly used to indicate human-generated responses is sometimes translated to other languages by mail user agents, or otherwise interpreted by mail user agents as indication that the message is a reply, so the

(Greek) prefix "Auto:" may also be translated or used as a generic indication that the message is an automatic response. However the "Auto:" indication is intended only as an aid to humans in processing the message. Mail processing software SHOULD NOT assume that the presence of "Auto:" at the beginning of a Subject field is an indication that the message was automatically submitted.

Note that the Subject field of the subject message may contain encoded-words formatted according to [N3] and [n3.5], and such text MAY be included in the Subject field of a response. In generating responses containing such fields there is rarely a need to decode and re-encode such text. It is usually sufficient to leave those encoded-words as they were in the subject message, merely prepending "Auto: " or other indication. However, it is still necessary to ensure that no line in the resulting Subject field that contains an encoded-word is greater than 76 ASCII characters in length (this refers to the encoded form, not the number of characters in the text being encoded). Also, if the responder truncates the Subject from the subject message it is necessary to avoid truncating Subject text in the middle of an encoded-word.

3.1.6 In-Reply-To and References fields

The In-Reply-To and References fields SHOULD be provided in the header of a response message if there was a Message-ID field in the subject message, according to the rules in [N2] section 3.6.4.

3.1.7 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.

3.1.8 Precedence field

A response MAY include a Precedence field [I4] in order to discourage responses from some kinds of responders which predate this specification. The field-body of the Precedence field MAY consist of the text "junk", "list", "bulk", or other text deemed appropriate by the responder. Because the Precedence field is non-standard and its interpretation varies widely, the use of Precedence is not specifically recommended by this specification, nor does this specification recommend any particular value for that field.

3.2 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, especially if in doing so it is desirable to use multiple charsets.

Response messages SHOULD NOT include significant content from the subject message. In particular, Personal and Group responses SHOULD NOT contain non-text content from the subject message, and they SHOULD NOT include attachments from the subject message. Neither of these conditions applies to responders that specifically exist for the

purpose of altering or translating content sent to them (for instance, a FORTRAN-to-C translator); however, such responders MUST employ measures to avoid being used as a means of laundering or forwarding undesirable content, such as spam or viruses.

Note that when text from the Subject or other fields from the header of the subject message is included in the body of the response, it is necessary to decode any encoded-words that appeared in those fields before including in the message body, and to use an appropriate content-type, charset, and content-transfer-encoding. In some cases it may be necessary to transliterate text from the charset(s) used in the header of the subject message, to the charset(s) used in the body of the response. (It is much easier to implement a responder if text from the header of the subject message never needs to appear in the body of the response.)

3.2.1 Use of DSNs and MDNs instead of this specification

In general, it is appropriate to use Delivery Status Notifications (DSNs) for responses that are generated by the mail transport system as a result of attempts to relay, forward, or deliver mail, and only when the purpose of that response is to provide the sender of the subject message with information about the status of that mail delivery. For instance, a "virus scanner" which is activated by a mail delivery process to filter harmful content prior to delivery, could return a DSN with the Action field set to "failed" with a Status code of 5.7.1 (Delivery not authorized, message refused) if the entire message was not delivered due to security reasons; or it could return a DSN with the Action field set to "relayed" or "delivered" (as appropriate) with a Status code set to 2.6.4 (conversion with loss performed) if the message was relayed or delivered with the presumably harmful content removed. The DSN specification [I2], rather than this document, governs the generation and format of DSNs.

Similarly, it is appropriate to use Message Disposition Notifications (MDNs) only for responses generated on the recipient's behalf, which are generated on or after delivery to a recipient's mailbox, and for which the purpose of the response is to indicate the disposition of the message. The MDN specification [I3], rather than this document, governs the generation and format of MDNs.

This document is not intended to alter either the DSN or MDN specifications. Responses that fit within the criteria of DSN or MDN, as defined by the respective specifications, should be generated according to the DSN or MDN specification rather than this document. Responses which do not fit one of these sets of criteria should be generated according to this document.

3.3 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 will (of course) 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 [N5].

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

In general, automatic responses SHOULD be sent to the Return-Path field if generated after delivery. If the response is generated prior to delivery, the response SHOULD be sent to the reverse-path from the SMTP MAIL FROM command, or (in a non-SMTP system) to the envelope return address which serves as the destination for nondelivery reports.

If the response is to be generated after delivery, and there is no Return-Path field in the subject message, there is an implementation error 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. Such heuristics have been known to cause problems in the past.

A Service Responder MAY deliver the response to the address(es) from the From field, or to another address from the request payload, provided this behavior is precisely defined in the specification for that service. Services responders SHOULD NOT use the Reply-To field 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. For instance, a human sender may use Reply-To to request that replies be sent to an entire mailing list. 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 lacks 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. For these reasons the From field SHOULD NOT be used as a destination for automatic responses.

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, using the ABNF notation of [N6]:

```
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] parameter )
```

The symbols "token", and "parameter" are as defined in [N7] (as amended by [N4]).

The maximum number of Auto-Submitted fields that may appear in a message header is 1.

5.2 Semantics

The Auto-Submitted header field SHOULD NOT be supplied for messages that were manually submitted by a human. (However, user agents that allow senders to specify arbitrary fields SHOULD NOT prevent humans from setting the Auto-Submitted field,

because it is sometimes useful for testing.)

The auto-generated keyword:

- SHOULD be used on messages generated by automatic (often periodic) processes (such as UNIX "cron jobs") which are not direct responses to other messages,
- 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, except for messages which are automatic responses to other messages.

The "no" keyword MAY be used to explicitly indicate that a message was originated by a human, if for some reason this is found to be appropriate.

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 [N8]. Keywords beginning with "x-" are reserved for experiments and use among consenting parties. Recipients of messages containing an Auto-Submitted field with any keyword other than "no" MAY assume that the message was not manually submitted by a human.

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 from [N2] can be used to indicate a reason why this message was automatically submitted.

6. Security Considerations

Automatic responders introduce the potential 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 mode", thus taxing the resources of the mail transport system;

- Use of such responders to determine whether recipient addresses are valid, especially when such information is not otherwise provided (e.g. SMTP RCPT or VRFY command responses) and is not intended to be disclosed;
- Use of such responders to obtain personal information about recipients, including information about recipients' recent usage of his mailbox or recent activity;
- 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 use of responders to launder or amplify attacker-chosen 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 mode),
- Recommending that responses from Personal and Group Responders 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, any network-based service (including those accessed by email) needs to provide security that is sufficient to prevent the service from being used as a means to inappropriately or destructively access the resources that are accessible by the service.

It has also been noted that Personal and Group Responders sometimes inappropriately disclose recipients' personal information. This might happen automatically (as when a Group Responder automatically supplies a recipient's personal or mobile telephone number as alternate contact information) or "manually". Automatically-generated information SHOULD NOT include personal information about the recipient which is not

already known to, or easily available to, the sender of the subject message. User interfaces which allow recipients to supply response text SHOULD make it clear to the user that this information will be made available not only to local colleagues but also to the entire Internet, including potential attackers.

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 parameters are approved (through an IETF Consensus process) it may be appropriate for IANA to create a registry of such keywords or parameters.

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" [I7]. 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 [I8]. [I9] defined an "Autosubmitted" field for use when gatewaying between X.400 and Internet mail. Jacob Palme wrote an internet-draft [I10] 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: Russ Allbery, Adam Costello, Ned Freed, Lawrence Greenfield, Arnt Gulbrandsen, Eric Hall, Tony Hansen, Dan Kohn, Bruce Lilly, der Mouse, Lyndon Nerenberg, 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. Normative References

[N1] Bradner, S. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119, March 1997.

[N2] Resnick, P. (ed.) Internet Message Format. RFC 2822, April 2001.

[N3] Moore, K. MIME (Multipurpose Internet Mail Extensions) Part Three: Message Header Extensions for Non-ASCII Text. RFC 2047, November 1996.

[N4] Freed, N., Moore., K. MIME Parameter Value and Encoded Word Extensions: Character Sets, Languages, and Continuations. RFC 2231, November 1997.

[N5] Moore, K. SMTP Service Extension for Delivery Status Notifications. RFC 3461, January 2003.

[N6] Crocker, D. (ed.), Overell, P. Augmented BNF for Syntax Specifications: ABNF. RFC 2234, November 1997.

[N7] Freed, N. Borenstein, N. Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies. RFC 2045, November 1996.

[N8] Narten, T., Alvestrand, H. Guidelines for Writing an IANA Considerations Section in RFCs. RFC 2434, October 1998.

11. Informative References

[I1] "Sorcerer's apprentice mode", originally from the Jargon file once maintained at MIT-AI and SAIL; now collected at various places on the net. See e.g. `http://www.jargon.net/`

[I2] Moore, K. Vaudreuil, G. An Extensible Message Format for Delivery Status Notifications. RFC 3464, January 2003.

[I3] Fajman, R. An Extensible Message Format for Message Disposition Notifications. RFC 2298, March 1998.

[I4] Palme, J. Common Internet Message Headers. RFC 2076, February 1997.

[I5] Neufeld, G., Baer, J. The Use of URLs as Meta-Syntax for Core Mail List Commands and their Transport through Message Header Fields. RFC 2369, July 1998.

[I6] Crocker, D. Standard for the format of ARPA Internet text messages. RFC 822, August 1982.

[I7] 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. (reference included only for attribution)

[I8] X.400. (perhaps someone can supply the correct reference for the first version of the X.400 document to define autosubmitted?)

[I9] Kille, S. MIXER (Mime Internet X.400 Enhanced Relay): Mapping between X.400 and RFC 822/MIME. RFC 2156, January 1998.

[I10] Palme, J. "The Auto-Submitted and Expires Headers in E-mail". Expired Internet-Draft "draft-ietf-mailext-new-fields-15.txt", February 1999. (reference included only for attribution)