ï»¿

Delivered-To Email Header Field

Abstract

   The address to which email is delivered might be different than any
   of the addresses shown in any of the content header fields that were
   created by the email's author. For example, the address used by the
   email transport service is provided separately, such as through
   SMTP's "RCPT TO" command, and might not match any address in the To:
   or cc: fields. In addition, before final delivery, handling can
   entail a sequence of submission/delivery events, using a sequence of
   different destination addresses that (eventually) lead to the
   recipient. As well, a receiving system's delivery process can
   produce local address transformations.

   It can be helpful for a message to have a common way to record each
   delivery in such a sequence, noting each address used in the sequence
   to that recipient, such as for (1) analyzing the path a message has
   taken, (2) loop detection, or (3) formulating the author's address in
   a reply message. This document defines a header field for this
   information.

   Email handling information discloses details about the email
   infrastructure, as well as about a particular recipient; this can
   raise privacy concerns.

   A header field such as this is not automatically assured of
   widespread use. Therefore, this document is being published as an
   Experimental RFC, looking for constituency and for operational
   utility. This document was produced through the Independent
   Submission Stream and was not subject to the IETF's approval process.

Status of This Memo

   This document is not an Internet Standards Track specification; it is
   published for examination, experimental implementation, and
   evaluation.

   This document defines an Experimental Protocol for the Internet
   community. This is a contribution to the RFC Series, independently
   of any other RFC stream. The RFC Editor has chosen to publish this
   document at its discretion and makes no statement about its value for
   implementation or deployment. Documents approved for publication by
   the RFC Editor are not candidates for any level of Internet Standard;
   see Section 2 of RFC 7841.

   Information about the current status of this document, any errata,
   and how to provide feedback on it may be obtained at
   https://www.rfc-editor.org/info/rfc9228.

Copyright Notice

Table of Contents

1. Introduction

   The address to which email is delivered might be different than any
   of the addresses shown in any of the content header fields
   [Mail-Fmt], such as the To: and cc: fields that were created by the
   author's Message User Agent (MUA) [Mail-Arch]. The address used by
   the Message Handling Service (MHS) is provided separately, in
   envelope information, such as through a "RCPT TO" command in [SMTP].

   As noted in Section 4.3.3 of [Mail-Arch], 'A transfer of
   responsibility from the MHS to a Recipient's environment (mailbox) is
   called "delivery".' That is, when the destination address is fully
   and successfully processed, and any additional processing is by an
   agent working on behalf of that address, the message has been
   delivered. Rather than placing the message into a recipient inbox or
   otherwise completing the handling of the message, that agent might
   create additional processing, including to one or more different
   addresses. Each transition of responsibility, from the MHS to an
   agent of a current addressee, constitutes a distinct delivery. Given
   handling sequences that can include aliasing, mailing lists, and the
   like, the transit of a message from its author to a final recipient
   might include a series of submission/delivery events. Also, the
   delivery process at a receiving system can produce local (internal)
   address transformations.

   Header fields that provide information about handling can be used
   when assessing email traffic issues and when diagnosing specific
   handling problems. To this end, it can be helpful for a message to
   have a common way to indicate each delivery in the handling sequence
   and to include each address that led to the final delivery. This can
   aid in the analysis of a message's transit handling.

   An additional use can be to aid in detecting a delivery sequence
   loop, based on a specific address. With a problematic loop, the same
   copy of a message is delivered to the same email address more than
   once. This is different from having different copies delivered to
   the same address, such as happens when a message is sent directly to
   an address, as well as via a mailing list. It is also different from
   having two copies of the same message arrive at the same, ultimate
   destination address, having been originally posted to two different
   addresses. Further, this is different from noting when a message
   simply transits the same Message Transfer Agent (MTA) more than once,
   which might be necessary, such as when it is processed through a
   mailing list; an MTA services many addresses.

   Delivering the same copy of a message more than once, to the same
   address, is almost certainly not an intended activity. An example of
   a problematic arrangement would be to send a message to mailing list
   List-A, where List-A contains an entry for List-B, and List-B
   contains an entry for List-A. The message will enter an infinite
   loop. Loop detection for email can be a complicated affair. The
   Delivered-To: header field provides helpful information, with a
   definitive indication that this copy of a message has (already) been
   delivered to a specific address.

When specifying new activity that is related to existing activity, there is a choice of design approach:

* Seeking to change (some of) the existing behavior

* Adding to the activity without changing what is already being done

* Calling for separate, new activity

On the average, attempting to change existing activities is the least likely to obtain adoption; it can create operational confusion between old and new activities, which in turn creates resistance to adoption. Seeking new activity can make sense when that activity is sufficiently different and deemed sufficiently beneficial. Adding to existing activity has the selling point of building upon an installed base. The current specification builds upon an existing installed base of Delivered-To: activity. It calls for little technical enhancement; rather, it simply provides for a wider range of application.

Considerations:

* Email handling information, such as this, provides information about the email infrastructure, as well as about the recipient. Disclosure of this information might engender privacy concerns.

* A specification is not automatically assured of adoption or use. Therefore, this document is being published as an Experimental RFC, looking for extended constituency and for general operational utility.

* This document was produced through the Independent RFC Stream and was not subject to the IETF's approval process.

2. Background

Ad hoc use of a Delivered-To: email header field appears to date back to the 1990s, primarily for loop detection, although documentation is spotty and system specific. A listing of some implementations is offered in [Prior].

It appears that all uses include a string in the form of an email address, although at least one example has leading text that is a comment about the address. In some cases, the string appears to be the email transport destination address, such as the address used in SMTP's "RCPT TO" command. In other cases, it appears to be the result of some internal mapping at the receiving system, although tending to be a variant of the transport address.

Email loop detection tends to be accomplished through a variety of different methods, such as counting Received: header fields. These methods are often combined to greater effect.

The Received: header field's 'for' clause is sometimes useful for disclosing the recipient's address. However, the clause is not used reliably, and its semantics are not thoroughly defined. Also, it references an addressing value that is received but might be different from the value that is ultimately used (as the result of a transformation). That is, the value in a 'for' clause might be a sufficient indicator of delivery addressing, but it might not.

3. Framework & Terminology

Unless otherwise indicated, basic architecture and terminology used in this document are taken from:

* [Mail-Arch]

* [SMTP]

*   [Mail-Fmt]

and syntax is specified with:

*   [ABNF]

Normative language is per [RFC8174]:

> The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL
> NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED",
> "MAY", and "OPTIONAL" in this document are to be interpreted as
> described in BCP 14 [RFC2119] [RFC8174] when, and only when, they
> appear in all capitals, as shown here.

4. Delivered-To

The Delivered-To: header field annotates an email delivery event.
The header field contains information about the individual address
used to effect that transition.

*   When a message is delivered, as a transition from control by the
    MHS to the recipient's store or their agent, a Delivered-To:
    header field SHOULD be added, with the _addr-spec_ value
    containing the address that was used by the service to reach the
    recipient.

*   If a receiving system's delivery process applies mappings or
    transformations from the address used by the MHS to a local value,
    this new value SHOULD also be recorded into a separate Delivered-
    To: field when transit and processing using that address
    successfully complete. This ensures a detailed record of the
    sequence of handling addresses used for the message.

*   As with some other information, each additional Delivered-To:
    header field MUST be placed at the current 'top' of the message's
    set of header fields -- that is, as the first header field, in a
    fashion similar to the trace fields specified in [SMTP] (for
    example, Section 4.1.1.4 of [SMTP]). This produces a sequence of
    Delivered-To: header fields that represent the sequence of
    deliveries, with the first being at the 'bottom' of the sequence
    and the final one being at the top.

*   As with other fields placed incrementally in this way, with each
    added at the current top, the Delivered-To: header field MUST NOT
    be reordered with respect to other Delivered-To: fields and those
    other fields. This is intended to preserve the fields as
    representing the message handling sequence.

The Delivered-To: header field is added at the time of delivery, when
responsibility for a message transitions from the Message Handling
Service (MHS) to an agent of the specified individual recipient
address. The field can also be added as a result of internal system
processing, to note address transformations.

> Note: The presence of an existing Delivered-To: header field,
> for the same address, typically indicates a handling loop for
> this instance of the message.

The syntax of the header field is:

"Delivered-To:" FWS addr-spec FWS CRLF ; addr-spec is from [Mail-Fmt]

The field records information about a single address, for one
recipient. See Section 6 for the privacy-related concerns about
divulging addresses.

5. Multi-Delivery Example

The Delivered-To: header field can be used to document a sequence of
deliveries of a message. Each time an address is fully processed, a
Delivered-To: header field is added, recording a handling sequence,

with the most recent one being towards the 'top' of the sequence of
header fields.

This example demonstrates a message traveling from its original
posting, through a remote group mailing list, on through an
independent personal aliasing mechanism, and then reaching final
delivery at yet another independent email provider.

1. Origination at com.example

        The message, as submitted. The destination address is the
        same as the value in the message's To: header field.

    From: Ann Author <aauthor@com.example>
    Date: Mon, 25 Jan 2021 18:29:06 -0500
    To: list@org.example
    Subject: [list] Sending through a list and alias
    Sender: Ann Author <aauthor@com.example>

2. List processing at org.example

        As delivered, with one Delivered-To: header field, to the list
        processing module, which will then resubmit the message for
        further transport to the list member "Recipient-
        alumn@edu.example".

 Delivered-To: list@org.example
 Received: by submit.org.example with SMTP id i17so174806891jn.1
     for <list@org.example> from mail.com.example;
     Mon, 25 Jan 2021 15:29:19 -0800 (PST)
 Received: by mail.com.example; Mon, 25 Jan 2021 15:29:00 -0800 (PST)
 From: Ann Author <aauthor@com.example>
 Date: Mon, 25 Jan 2021 18:29:06 -0500
 To: list@org.example
 Subject: [list] Sending through a list and alias
 Sender: Ann Author <aauthor@com.example>

3. Alias processing at edu.example

        The message, as delivered with two Delivered-To: header
        fields, to the alias processing module, which sends the
        message on to "theRecipient@example.net".

 Delivered-To: Recipient-alumn@edu.example
 Received: from mail.org.example
     by relay.edu.example; Mon, 25 Jan 2021 23:29:24 +0000 (UTC)
 Received: by submit.org.example;
     Mon, 25 Jan 2021 23:29:21 +0000 (UTC)
 Delivered-To: list@org.example
 Received: by submit.org.example with SMTP id i17so174806891jn.1
     for <list@org.example> from mail.com.example;
     Mon, 25 Jan 2021 15:29:19 -0800 (PST)
 Received: by mail.com.example; Mon, 25 Jan 2021 15:29:00 -0800 (PST)
 From: Ann Author <aauthor@com.example>
 Date: Mon, 25 Jan 2021 18:29:06 -0500
 To: list@org.example
 Subject: [list] Sending through a list and alias
 Sender: list-bounces@org.example

4. Final delivery to the recipient at example.net

        The message, as finally delivered with three Delivered-To:
        header fields, to the recipient at "theRecipient@example.net".

 Delivered-To: theRecipient@example.net
 Received: from mail.edu.example (mail.edu.example [4.31.198.45])
     by relay.example.net; Mon, 25 Jan 2021 23:29:24 +0000 (UTC)
 Delivered-To: Recipient-alumn@edu.example
 Received: from mail.org.example
     by relay.edu.example; Mon, 25 Jan 2021 23:29:24 +0000 (UTC)
 Received: by submit.org.example;

```
        Mon, 25 Jan 2021 23:29:21 +0000 (UTC)
    Delivered-To: list@org.example
    Received: by submit.org.example with SMTP id i17so17480689ljn.1
        for <list@org.example> from mail.com.example;
        Mon, 25 Jan 2021 15:29:19 -0800 (PST)
    Received: by mail.com.example; Mon, 25 Jan 2021 15:29:00 -0800 (PST)
    From: Ann Author <aauthor@com.example>
    Date: Mon, 25 Jan 2021 18:29:06 -0500
    To: list@org.example
    Subject: [list] Sending through a list and alias
    Sender: list-bounces@org.example
```

6. Security Considerations

    As with Received: header fields, the presence of a Delivered-To:
    header field discloses handling information and, possibly, personal
    information.

    Security and privacy are essential, if challenging, topics for email
    in general and for the handling of metadata in particular. The
    purpose of this section is to note points of potential concern,
    rather than to provide details for mitigation. The basic mechanism
    described here has a long history of use, with no history of being
    problematic. However, the expanded use described here might create
    new scenarios that are problematic.

    An issue specific to this mechanism is disclosure of a sequence of
    addresses, applied to the same recipient, if a message goes through a
    series of recipient address replacements. This document calls for
    each of these addresses to be recorded in a separate Delivered-To:
    field. This does not disclose addresses of other recipients, but it
    does disclose an address-transformation handling path for the
    recipient.

    This disclosure is most likely to be a concern when a recipient
    manually forwards a message and includes all of the original header
    fields. This will expose, to a later recipient, any intermediate
    addresses used for getting the original message to the original
    recipient. Such a disclosure is likely to be unintended and might be
    (highly) problematic. Note that a basic version of this unintended
    disclosure has long existed, by virtue of a later recipient's seeing
    Received: header fields, but especially any with a 'for' clause.
    However, a Delivered-To: header field sequence can disclose
    significantly more recipient-specific handling detail.

    An issue that is entirely implementation specific -- and therefore
    out of scope for this document -- is that in some systems, a message
    that is for multiple (local) recipients is stored as a single, shared
    version. Supporting Delivered-To:, while maintaining recipient
    privacy, creates a challenge in this case, since exposing different
    recipient addresses to other recipients can be problematic.

7. IANA Considerations

    IANA has registered the Delivered-To: header field as below, per
    [RFC3864] in the "Provisional Message Header Field Names" registry:

    Header Field Name: Delivered-To

    Protocol: mail

    Status: Provisional

    Author/Change controller: Dave Crocker

    Specification document(s): *** This document ***

    Related information: None.

8. Experimental Goals

Specific feedback is sought concerning:

*   Technical issues in recording the Delivered-To: field into a
    message, through its entire submission/delivery sequence

*   Market interest in the uses described here

*   Utility for the purposes described here, or for other uses

So the questions to answer for this Experimental RFC are:

*   Is there demonstrated interest by MSA/MTA/MDA (Message Submission
    Agent / Message Transfer Agent / Message Delivery Agent)
    developers?

*   If the capability is implemented and the header field generated,
    is it used by operators or MUAs?

*   Does the presence of the header field create any operational
    problems?

*   Does the presence of the header field demonstrate additional
    security issues?

*   What specific changes to the document are needed?

*   What other comments will aid in use of this mechanism?

Please send comments to ietf-smtp@ietf.org.

9. References

9.1. Normative References

[ABNF]   Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax
            Specifications: ABNF", STD 68, RFC 5234,
            DOI 10.17487/RFC5234, January 2008,
            <https://www.rfc-editor.org/info/rfc5234>.

[Mail-Arch]
            Crocker, D., "Internet Mail Architecture", RFC 5598,
            DOI 10.17487/RFC5598, July 2009,
            <https://www.rfc-editor.org/info/rfc5598>.

[Mail-Fmt] Resnick, P., Ed., "Internet Message Format", RFC 5322,
            DOI 10.17487/RFC5322, October 2008,
            <https://www.rfc-editor.org/info/rfc5322>.

[RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate
            Requirement Levels", BCP 14, RFC 2119,
            DOI 10.17487/RFC2119, March 1997,
            <https://www.rfc-editor.org/info/rfc2119>.

[RFC3864]   Klyne, G., Nottingham, M., and J. Mogul, "Registration
            Procedures for Message Header Fields", BCP 90, RFC 3864,
            DOI 10.17487/RFC3864, September 2004,
            <https://www.rfc-editor.org/info/rfc3864>.

[RFC8174]   Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC
            2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174,
            May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[SMTP]   Klensin, J., "Simple Mail Transfer Protocol", RFC 5321,
            DOI 10.17487/RFC5321, October 2008,
            <https://www.rfc-editor.org/info/rfc5321>.

9.2. Informative References

[Prior]   Dukhovni, V. and J. Levine, "The Delivered-To Message
            Header Field", Work in Progress, Internet-Draft, draft-
            duklev-deliveredto-01, 6 February 2022,

<https://datatracker.ietf.org/doc/html/draft-duklev-
            deliveredto-01>.

Acknowledgements

      Even a simple, narrow specification can elicit a remarkable range and
      intensity of debate. In spite of the current document's being a case
      of that challenge, useful discussion has taken place, first in the
      IETF's emailcore working group mailing list, and then on the long-
      standing ietf-smtp mailing list.

      Helpful information and suggestions were provided by Anonymous,
      StÃ©phane Bortzmeyer, Richard Clayton, Viktor Dukhovni, Adrian Farrel,
      Ned Freed, John Klensin, Barry Leiba, Brandon Long, George
      Michaelson, Michael Peddemors, Phil Pennock, Pete Resnick, Sam
      Varshavchik, Alessandro Vesely, and Tim Wicinski.

Author's Address

      Dave Crocker (editor)
      Brandenburg InternetWorking
      Email: dcrocker@bbiw.net