

Internet Draft
[draft-ietf-usefor-posted-mailed-01.txt](#)
Category-to-be: Proposed standard

Matt Curtin
The Ohio State University
Jamie Zawinski
Netscape Communications

July 1998
Expires: Six Months from above date

Identification of messages delivered via both mail and news

Status of this Memo

This document is an Internet-Draft. 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.''

To view the entire list of current Internet-Drafts, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Northern Europe), ftp.nis.garr.it (Southern Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), or ftp.isi.edu (US West Coast).

Abstract

This draft defines a format to be used when delivering a single message to multiple destinations, where some destinations are newsgroups and some destinations are email mailboxes.

Table of contents

- 1. Introduction**
- 2. Terminology**
- 3. Definitions of new message elements**
 - 3.1. The Posted-And-Mailed header**
 - 3.2. The Followup-Host header**
 - 3.3. The message body prolog**
- 4. Clarifications of the semantics of existing headers**
 - 4.1. References and In-Reply-To**
 - 4.2. Message-ID**
 - 4.3. Followup-To**
 - 4.4. Newsgroups**
- 5. Security considerations**
- 6. Acknowledgments**

[7. References](#)

[8. Authors' Addresses](#)

Appendix A: Examples

[1. Introduction](#)

Most news readers include facilities for generating replies which are either posted to news, or mailed directly to the author. An increasing number of news readers have the ability to do both simultaneously: to post one copy of the message to news, and to send a copy of that same message to a set of recipients via email.

When one receives an email message, there is currently no reliable way to identify that message as being one which has also been posted. This document specifies a mechanism by which such messages may be identified.

The model used in this document is that a single message is prepared, and then delivered on multiple transports to its various destinations. This is not intended to dictate anything about the mechanism by which message delivery must be implemented. But, it is rather intended to convey the intent that both messages should, as far as is possible, have identical bodies and headers.

Obviously, various transports (including, but not limited to, [\[SMTP\]](#) and [\[NNTP\]](#)) will add various headers to the messages they carry, and so it will never be the case that two copies of the same message which are received via different transports will be identical. However, excepting the headers added by the transports, it is likely that the two copies of the message will have identical headers, and is also likely that they will have identical bodies.

It is also recognized that certain elements of the transport (including, but not limited to, mail-news gateways, mailing list reflectors, and newsgroup or mailing list moderators) might modify existing message bodies and headers. The modification might be in form only, such as the Content-Transfer-Encoding ([\[MIME\]](#)) of the body being changed; or it might be substantive, such as a standard disclaimer, or standard set of instructions, being appended to the bodies. This means that software conforming to this document cannot guarantee that the two messages will have identical bodies by the time they reach their destinations. It can, however, hold that as a goal, with the recognition that the goal will not always be reachable due to forces beyond its control. (In other words, the authors believe that transport and gateway software that so alters the message bodies is generally wrong in so doing, while recognizing that sometimes it's a choice between evils and imperfect preservation of the message may be the best that can be done.)

[2. Terminology](#)

In this document, we will discuss several "views" on the same logical message.

A "combined message" is a message for which the sender has chosen some email destinations, as well as and some news destinations. When we are talking about a combined message, we are either talking about it as it appeared when composed, but before it was delivered; or we are talking about both its mail and news aspects simultaneously.

When this document refers to a "mail message", it refers to a message received by someone via email. In isolation, "mail message" could refer to either a simple mail-only message, or to the version of a combined message that was delivered by mail.

Likewise, unless otherwise specified, "news message" refers to a message that was received by someone via news, and that may or may not also have been mailed.

This document assumes familiarity with the mail and news message formats, as documented in [[MAIL](#)] and [[NEWS](#)].

3. Definitions of new message elements

3.1. The Posted-And-Mailed header

When a message is sent to both mail and news destinations, it MUST include a Posted-And-Mailed header, with the value "yes":

```
posted-and-mailed = "Posted-And-Mailed:" [CFWS] ("yes" / "no") [CFWS] CRLF
```

The presence of this header indicates that the message conforms to this document. If this header is absent, conformant software MUST assume that the message in question DOES NOT conform to this document, and MUST NOT treat the message as a combined message. If this header is present but has a value other than "yes", conformant software MUST NOT treat the message as a combined message.

This header MUST be present both in the version of the message which was sent to a news transport, and in the version of the message which was sent to a mail transport, and MUST be identical in both.

This header MAY be omitted in messages which are not combined messages. This is a reasonable thing to do, because messages which do not have that header MUST be assumed by receiving agents to be non-combined messages. However, if this header is included in a non-combined message, it MUST have the value "no".

3.2. The Followup-Host header

This header is optional.

If it is present, it is an instruction to the recipient about what news host and protocol SHOULD be used to send a reply, should the recipient desire to send a reply to any of the newsgroups listed in the Followup-To or Newsgroups headers.

Background:

It is becoming more common for discussion forums to exist which are for all practical purposes newsgroups, but which are served by only one (or a small number of) hosts. They are not widely replicated.

The way one uses these groups is by connecting to a particular port on a particular host and speaking a particular news protocol (typically NNTP.)

This differs from the traditional USENET model, where one connects to a local news server for all activity, and the messages are propagated to many different hosts.

It is not the place of this document to discuss the pros or cons of this mode of operation. However, this document recognizes that this mode of operation exists, and defines a mechanism to deal with the issues related to posted-and-mailed messages as relates these non-USENET news hosts, as well as the more common USENET case. It is noteworthy that other factors (such as Internet gateway architectures which prohibit connectivity between "internal" clients and "external" news hosts) might prevent the mechanism from working as desired.

The Followup-Host header SHOULD be used when all newsgroups in the Newsgroups and Followup-To headers are served by a single, non-USENET news server.

It MUST NOT be used when the newsgroups in question use the traditional USENET model of propagation: that is, newsgroups which are not ones that are served only by a particular host.

The Followup-Host header is an instruction to use a particular host for posting activity. Therefore, its use includes the assumption that the recipients of the message will be able to post via the host in question.

It is recognized that even traditional USENET groups have varying levels of propagation, and that there is no guarantee that any mail recipient has access to any server which offers a particular USENET group. The Followup-Host header is not intended to address this problem.

How the posting software makes the determination of whether the current news server is a USENET-style server, or a non-USENET style server is unspecified. It is left up to the implementor. If the client cannot make that determination, then the client MUST assume that the newsgroup is a USENET-style one, and therefore MUST NOT include a Followup-Host header.

One possible way to make the determination would be for software which was able to deal with multiple news hosts to remember which hosts were USENET and which were not. A particular news agent might have a notion of a "default" host, and assume that the default host was USENET, and the non-default hosts were not. Another news agent might ask the user to specify whether the host carried USENET at the time the user connected to the host (or subscribed to a group carried by it.)

The body of the Followup-Host header is a URL, as defined by [[NEWSURL](#)]:

```
followup-host = "Followup-Host:" [CFWS] news-url [CFWS] CRLF
```

```
news-url = <a URL representing a news service>
```

The reason for providing a full URL rather than simply a host name is that news service may not necessarily be provided by [[NNTP](#)]. URLs, being extensible, provide an easy way to accommodate current and future protocol innovations.

The header's contents could be as simple as:

```
Followup-Host: news://news.example.com/
```

indicating the default news protocol (NNTP) on the default nntp port (TCP 119). An NNTP service running on a nonstandard port could be expressed as

```
Followup-Host: news://news.example.com:6666/
```

A news service running a protocol other than NNTP would be expressed by using a different type of URL. For example, this header represents news service running on the "snews" protocol (which is actually NNTP wrapped inside of SSL):

```
Followup-Host: snews://secnews.netscape.com/
```

It is beyond the scope of this document to document these protocols or URL syntaxes.

[3.3. The message body prolog](#)

When a message is sent to both mail and news recipients, the posting software MAY choose to automatically include a free-text blurb at beginning of the message body indicating that it has been posted as well as mailed.

If this text is inserted, it MUST be inserted in BOTH the copy of the message that is posted, and in the copy that is mailed. This is in keeping with the guiding principle that two copies of the same message MUST have the same Message-ID, and that, conversely, two messages with

the same Message-ID MUST have the same body.

Message reading software MUST NOT attempt to automatically parse or otherwise interpret this body text. Such software should use the appropriate message headers instead. This body text, like all body text, is intended only for human consumption.

If the text is inserted, it SHOULD be kept brief: it is recommended that it consist only of one or two lines of text. For example,

```
[[ This message was both posted and mailed.  ]]
```

or perhaps

```
[[ This message was both posted and mailed: see  
   the `To' and `Newsgroups' headers for details.  ]]
```

But note that more verbose prologs are allowed, if desired by the user and/or the user's software.

Rationale:

The reason that a user or user-agent might want to insert a body prolog at all is to draw attention to the fact that this is a combined message. Historically, mail readers did not show the Newsgroups header by default, and news readers did not show the To and CC headers by default; therefore, it is likely that, in the absence of the body prolog, a user might mistakenly assume that a combined message was mailed only, or posted only.

The body prolog, if present, is largely redundant with the message headers. This redundancy is called for due to their different uses: the Posted-And-Mailed header is for interpretation by programs; the body is for interpretation by humans. It is intended that when support for the Posted-And-Mailed header becomes more widespread in mail and news reading software, the use of the body prolog will become unnecessary, and deprecated.

There are two reasons that the same text MUST be present in both the mailed and posted copies, if it is present at all.

- * The first reason is that it is a part of the message body; and having two different messages, with different bodies, but with the same Message-ID, would be a misuse of the Message-ID header. Therefore, if it was desired that the bodies differ, then one might conclude that the two different messages should have different Message-IDs. However, it is HIGHLY desirable for the two messages to have the same ID (as discussed in [section 4.2](#), "Message-ID".) Therefore, since the two messages MUST have the same ID, they MUST have identical bodies.

Some might argue that the bodies are "substantially the same", and that perhaps an exception should be made, and the two

messages with non-identical bodies should be allowed to have the same Message-ID anyway. The problem with this is that it is a slippery slope: it sets a bad precedent. Where would it end? Should it be allowed for two messages to have the same ID if one of them is in plain-text and the other is in HTML? If one includes alternate forms of attached documents? If one has been spell-checked and the other has not? If one is in English, and the other is a translation into Spanish? And so on.

- * The second reason is that the body prolog provides useful information to all recipients, regardless of whether they receive the message via news or mail. When one sends a mail message To: Bob, and CC: Alice, one does not send Bob a message that contains only a To field, and Alice a message that contains only a CC field. Rather, one discloses the full set of recipients to all recipients.

The rationale for having a body prolog at all is the assumption that the message headers (To, CC, and Newsgroups) are not sufficient to fully disclose the set of recipients of the message, because readers will tend not to be shown those headers by default.

If one accepts that the body prolog is necessary for full disclosure of the set of recipients to a mail recipient, then one must also accept that it is necessary for full disclosure to a news recipient.

There is controversy over whether it is appropriate for a user agent to insert this text at all, the argument against it being that any attempt to impose structure on message bodies is both inappropriate, and doomed to fail.

However, it is already the case that user agents are free to insert, automatically or otherwise, any manner of text into message bodies: for example, signature files, or "message templates." The reason this document mentions the possibility of a body prolog which labels combined messages is simply to ensure that conformant software which DOES choose to insert such a blurb does so in both copies of the message, not just one (for the reasons enumerated above.)

4. Clarifications of the semantics of existing headers

The general principle used here is that when a header is required in either mail or news, a combined message should include both headers. Combined with the principle that the same message text be delivered to both transports, this means that certain previously-news-only headers will be delivered over mail transports, and certain previously-mail-only headers will be delivered over news transports.

When sending a message as both mail and news, that message MUST meet the underlying requirements of both mail messages and news messages simultaneously.

[4.1.](#) References and In-Reply-To

Messages which are delivered to both mail and news MUST use the news [[NEWS](#)] syntax and semantics of the References header, since that RFC has more restrictive (and, arguably, more useful) syntax and semantics than does the mail message standard [[MAIL](#)].

Messages which are delivered to both mail and news, and which are replies, MUST have a References header.

If the Author-Message-ID header is present, messages sent in reply MUST include it in the References header.

Messages which are delivered to both mail and news MAY include an In-Reply-To header, with the semantics defined in [[MAIL](#)]. Should an In-Reply-To header be used, it MUST contain the last message ID of the References header (that is, the ID of the message to which this is a reply.)

[4.2.](#) Message-ID

Messages which are delivered to both mail and news MUST have identical Message-ID headers.

The syntax of the Message-ID header MUST be as defined in [[NEWS](#)], as that is a more restrictive subset of the syntax defined in [[MAIL](#)].

The Message-ID header is optional in [[SMTP](#)]. Generally, if the user agent does not generate the Message-ID, then the transport will generate one for the message. (This is typically true in the case of both news and mail. It is noteworthy that this is not a requirement for news hosts. Those that will behave this way go beyond the specification.)

Since allowing the server(s) to generate the IDs would cause the use of two different Message-IDs, in order to comply with this rule, a client will probably need to generate the Message-ID before handing the message to either transport. Some suggestions for good client-side Message-ID generation are offered in [[MSGID](#)].

(It is conceivable that some future message submission protocol might allow the client to ask the server to generate and return a Message-ID for it, but this is not possible with any of the currently-existing message submission protocols. So, the requirement is that the two copies of the message MUST have identical Message-IDs, but any mechanism which achieves this end is acceptable.)

Rationale:

The requirement that the Message-IDs be identical is to make it possible for a recipient of a combined message to reply to it and generate a correct, usable References header.

For example: if a Bob sends a combined message to a newsgroup and CCs the message to Alice by mail, Alice might want to reply in public, rather than in private: she might want to post her reply to the same newsgroup. If Bob uses the same Message-ID on both, that works. But, if Bob uses two different message IDs, then the message ID in Alice's References header will be different depending on whether she replies to the mailed copy or the posted copy. If she replies to the mailed copy, her new news message will mention an ID in its References field which is not present in the newsgroup. Therefore, the news thread will be fractured.

Similar fracturing effects can occur in mail, when combined messages are delivered to multiple mail recipients, or to newsgroups.

[4.3.](#) Followup-To

If both Posted-And-Mailed (with value "yes") and Followup-To are present, then replies which are to be posted MUST be directed to the newsgroups listed in the Followup-To header by default.

If a Followup-To header is present but a Posted-And-Mailed header is not (or has a value other than "yes"), then:

- * For a news message, the proper interpretation is defined by [[NEWS](#)].
- * For a mail message, the Followup-To header MUST be ignored.

[4.4.](#) Newsgroups

If the Posted-And-Mailed header is present and has the value "yes", then the Newsgroups header MUST also be present. In that case, the Newsgroups header has the semantics defined by [[NEWS](#)]: it lists the newsgroups to which this message was posted. This is true regardless of whether the message in question is the mailed copy or the posted copy: the Newsgroups header has the same semantics in both copies.

If a Newsgroups header is present but the Posted-And-Mailed header is not, or if the Posted-And-Mailed header has some value other than "yes":

- * For a news message, the proper interpretation is defined by [[NEWS](#)].
- * For a mail message, the Newsgroups header MUST be ignored.

Rationale:

The requirement to ignore lone Newsgroups headers in mail messages is an important one. Existing practice does not allow one to make any assumptions about the interpretation of the Newsgroups header in mail, as there are two widely used, conflicting interpretations of it: some message-generating software uses it as an indication that this mail message was also posted (for example, Gnus, Pine, and Netscape); and some message-generating software uses it as an indication of the groups to which the message to which this message is a reply was posted (for example, rn and its direct descendants like trn and slrn.)

Therefore, one MUST NOT interpret the Newsgroups header in a mail message unless that message is known to be in conformance with this document: unless there is a Posted-And-Mailed header in the message, the semantics of the Newsgroups header is undefined (and unsafe.) If there is a Posted-And-Mailed header, then the semantics of the Newsgroups header IS defined: by this document.

5. Security considerations

This format will reduce the risk of various unexpected results for combined messages. Some existing risks in email and news may stay even with this format, but no new risks are expected as a result of using this format. In general, increased transportation of messages between news and email may mean that existing risks in news are propagated to email or the reverse, but these risks would not be reduced by the lack of a standard for such combined messages.

The union of the security and privacy risks of existing mail and news usage must be considered; for example, care should be taken not to inappropriately disclose the BCC recipients of a mailed message to the news recipients.

6. Acknowledgments

This document is derived from and inspired by earlier proposals written by Jacob Palme and John Stanley. Valuable feedback was provided by Terje Bless, Stefan Haller, Lars Magne Ingebrigtsen, John Moreno, Jacob Palme, Pete Resnick, Bart Schaefer, Jeroen Scheerder, Brad Templeton, and Curtis Whalen.

7. References

Ref.	Author, title	IETF status (June 1998)
---	-----	-----
[SMTP]	J. Postel: "Simple Mail Transfer Protocol", STD 10, RFC 821 , August	Standard, Recommended.

1982.

[MAIL]	D. Crocker: "Standard for the format of ARPA Internet text messages." STD 11, RFC 822 , August 1982.	Standard, Recommended.
[NNTP]	B. Kantor, P. Lapsley: "Network News Transfer Protocol", RFC 977 , February 1986.	Non-standard (but still widely used as a de-facto standard).
[NEWS]	M.R. Horton, R. Adams: "Standard for interchange of USENET messages", RFC 1036 , December 1987.	Non-standard (but still widely used as a de-facto standard).
[MIME]	N. Freed, N. Borenstein and others, "Multipurpose Internet Mail Extensions (MIME) Part One to Five", RFC 2045 to 2049.	Draft Standard, elective.
[NEWSURL]	T. Berners-Lee, L. Masinter and others, "Uniform Resource Locators", RFC 1738 .	Draft Standard.
	See also: A. Gilman, "The 'news' URL scheme", draft-gilman-news-url-01.txt .	Internet Draft, work in progress.
[USEFOR]	D. Ritter, "News Article Format", draft-ietf-usefor-article-01.txt .	Internet Draft, work in progress.
[MSGID]	M. Curtin and J. Zawinski, "Recommendations for generating Message-IDs", draft-ietf-usefor-message-id-01.txt .	Internet Draft, work in progress

8. Authors' Addresses

Matt Curtin
The Ohio State University
791 Dreese Laboratories
2015 Neil Ave
Columbus OH 43210
+1 614 292 7352
cmcurtin@cis.ohio-state.edu

Jamie Zawinski
Netscape Communications Corporation
501 East Middlefield Road
Mountain View, CA 94043

(650) 937-2620
jwz@netscape.com

Appendix A: Examples

The following is an example of a combined message, sent both to a newsgroup comp.lang.c and via e-mail to a person bob@foo.example.com.

Here is the message as prepared by the message composition software:

```
Date: 7 Jan 1997 12:34:21 +0000 (GMT)
From: alice@bar.example.com
Subject: language or grade
Message-ID: <123zx@example.com>
To: bob@foo.example.com
Newsgroups: comp.lang.c
Posted-And-Mailed: yes
```

Which is it?

The same message as it might be received by someone reading it in the newsgroup comp.lang.c:

```
Path: news1.example.com!news2.example.com!bar.example.com!alice
NNTP-Posting-Host: news2.example.com
Xref: news.blatt.example.com comp.lang.c:20465
Lines: 1
Date: 7 Jan 1997 12:34:21 +0000 (GMT)
From: alice@bar.example.com
Subject: language or grade
Message-ID: <123zx@example.com>
To: bob@foo.example.com
Newsgroups: comp.lang.c
Posted-And-Mailed: yes
```

Which is it?

The same message as it might be received by a mail recipient (presumably bob@foo.example.com):

```
Return-Path: <alice@bar.example.com>
Received: from foo.example.com [127.0.0.1] by quux.example.com
Received: from quux.example.com [127.0.0.1] by bar.example.com
Date: 7 Jan 1997 12:34:21 +0000 (GMT)
From: alice@bar.example.com
Subject: language or grade
Message-ID: <123zx@example.com>
To: bob@foo.example.com
Newsgroups: comp.lang.c
Posted-And-Mailed: yes
```

Which is it?