

Internet Engineering Task Force (IETF)
Request for Comments: 6729
Category: Standards Track
ISSN: 2070-1721

D. Crocker
Brandenburg InternetWorking
M. Kucherawy
Cloudmark, Inc.
September 2012

Indicating Email Handling States in Trace Fields

Abstract

This document registers a trace field clause for use in indicating transitions between handling queues or processing states, including enacting inter- and intra-host message transitions. This might include message quarantining, mailing list moderation, timed delivery, queuing for further analysis, content conversion, or other similar causes, as well as optionally identifying normal handling queues.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6729>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Key Words	3
3. Trace Clause	3
4. Discussion	6
5. Granularity	6
6. IANA Considerations	7
6.1. MAIL Parameters Additional-registered-clauses Sub-Registry	7
6.2. MAIL Parameters Registered-states Sub-Registry	7
7. Security Considerations	9
8. References	10
8.1. Normative References	10
8.2. Informative References	10
Appendix A. Trace Field Examples	11
A.1. Typical Delivery without Obvious Extra Handling	11
A.2. Delivery with Moderation	11
Appendix B. Acknowledgements	12

1. Introduction

[SMTP] defines the content of email message trace fields, commonly the "Received" header field. These are typically used to record an audit trail of the path a message follows from origin to destination, with one such field added each time a message moves from one host to the next.

Section 3.7.2 of that document mentions that "the most important use of Received: lines is for debugging mail faults [...]".

There are some cases where there may be large time gaps between trace fields. Though this might be caused by transient communication issues, they might also be caused by policy decisions or special processing regarding the content of the message, authorization of some identity on the message, or transitions between major software components. Common examples include message quarantines (filters that cause a message to be held pending further evaluation or delivery of a message pending manual operator action), pending content analysis, or mailing list servers that impose moderation rules (mailing list owner action required regarding mail from authors not subscribed to those lists).

This document registers a new optional clause that can be used in trace fields to indicate that a message entered such a special processing queue or state for some period. This allows inspection of the trace information to reveal that the cause for a time gap in trace fields was imposed by additional processing rather than one caused by transient technical difficulties.

2. Key Words

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 [KEYWORDS].

3. Trace Clause

This specification defines a clause, called "state", which MAY be used when creating a Received header field (see Section 4.4 of [SMTP]) to indicate the nature of additional handling imposed on the relaying of a message toward its recipient(s). It is followed by a single keyword that provides that detail. A Mail Transfer Agent (MTA) or other handling agent that determines a message has entered a state other than normal queuing of messages for relaying or delivery MAY generate a trace field including one of these clauses. That is, the presence of this clause on a trace field is an indication of the entry of the message into that state; a later trace field added would indicate its departure from that state.

An MTA implementing this specification SHOULD add a Received field as described whenever:

- a. It determines that a special handling condition will occur and places it into that condition; or
- b. It determines that no special handling is required and prepares it for relay to the next handling agent.

An MTA need not add a Received field indicating preparation for normal handoff to the next handling agent if it has already added a Received field for some other reason. Trace data added by the next handling agent will imply the message's exit from the special handling condition.

If a single MTA processes a message through multiple special handling conditions, it MAY add a Received field for each distinct condition.

For example, presume a message will be injected into MTA-1, then travel to MTA-3 via MTA-2, and then MTA-3 will enact final delivery. At MTA-2, it is determined that some action will be taken that will cause the message to undergo some handling change that is outside of typical message flow. In this case:

1. MTA-1 adds a typical Received field and relays it to MTA-2.
2. MTA-2 determines that the atypical handling will occur and adds a Received field using the extension specified here.
3. On completion of the atypical handling, MTA-2 relays the message to MTA-3.
4. MTA-3 adds a typical Received field and enacts final delivery of the message.

Appropriate use of this mechanism does not include associating meta-data with the message, such as categorizing the message (e.g., the notions of "is spam" or "was 8-bit, converted to 7-bit"). Processing agents also cannot reliably use this mechanism to determine anything about the message content, since there is no guarantee that all agents in the chain of handling made such annotations to allow correct conclusions. The sole purpose here is to allow one to determine the point(s) in the chain of custody of a message at which the message was subjected to handling outside of normal message routing and queuing.

The following state keywords are defined in this document; extensions may define other registered keywords (see Section 6.2):

auth: The message entered a queue pending authentication of some identifier in the message.

content: The message entered a queue pending content analysis, such as scanning for spam or viruses.

convert: The message entered a queue pending content conversion.

moderation: The message entered a hold pending mailing list moderator action.

normal: The message is not in an administrative hold and is queued for or is being handed off to the next handling agent (which may be local delivery). This is the default interpretation when no "state" clause is present.

other: The message entered a hold or queue for reasons not covered by other keywords in this list and not for transient technology issues.

outbound: The message entered a queue for outbound relaying. This is typically the last case added for a single host, and the next Received header field is expected to be added by some other host.

quarantine: The message entered a hold in an isolation queue pending operator action for local policy reasons.

timed: The message entered a hold in order to meet a requested delivery window, such as is defined in [FUTURERELEASE].

In Section 6, the "state" clause is added to the Additional-Registered-Clauses IANA sub-registry. The ABNF for this clause is:

```
State = CFWS "state" FWS queue-state-keyword [ "/" value ]
queue-state-keyword = ( reg-state-keyword / unreg-state-keyword )
reg-state-keyword = ( "auth" / "content" / "convert" /
    "moderation" / "normal" / "other" /
    "outbound" / "quarantine" / "timed" /
    additional-state-keyword )
additional-state-keyword = token
    ; MUST be registered; see
    ; "IANA Considerations" below
value = token
unreg-state-keyword = token
```

"FWS" and "CFWS" are defined in [MAIL]. "token" is defined in [MIME].

A transfer agent making use of this extension MAY also include header field comments to provide additional information.

The "value" is available for providing additional labels as explanation for the state transition. Examples could include:

- o convert/unicode2ascii
- o moderation/not-subscribed
- o quarantine/spam

4. Discussion

Handling agents are not expected to implement or support all of these. Indeed, recording trace information for all of the states described above could make the header of a message inordinately large. Rather, an agent is encouraged to apply state annotations when a message enters a handling queue where a significant condition occurs or where significant additional processing or delay is possible, and especially when a handoff has occurred between two different, independent agents.

For example, an MTA receiving a message, doing message authentication, scanning for viruses and spam, and then putting it in an outbound queue could add four Received header fields denoting each of these states. However, where they are all done as part of a single system process, in a single pass, doing so would be considered unusual (and extremely verbose). This method SHOULD NOT be applied except when doing detailed analysis of a single component to identify performance issues with those steps.

Rather, an agent that wishes to make a state annotation SHOULD add only a single Received header field including such annotation, thus indicating (a) the time of completion of its handling of the message via the date portion of the field and (b) the final disposition of that message relative to that agent. For example, an MTA receiving a message that performs various checks on the message before immediately handing it off to a Mailing List Manager (MLM) would only record a "normal" state, assuming it passes those checks. The MLM would then evaluate the message and record its own state once it decides what the next step will be for the handling of that message.

5. Granularity

The degree of granularity -- and therefore the degree of verbosity -- recorded through the use of this additional trace clause is likely to vary depending on circumstances. It will typically be the case that use of this clause will be limited to "unusual" transitions, such as when a message requires additional scrutiny or other processing or needs to be quarantined.

Somewhat greater granularity might also include transitions of administrative responsibility, such as between a Mail Transfer Agent (MTA) operator and a Mailing List Manager (MLM) operator. This could be further enhanced to note some transitions that are interesting only when other transitions have occurred, such as noting entry to the outbound queue only when the message is originating from an "interesting" source, like an MLM, since an MLM can introduce significant changes to the message or delivery delay and it could be

useful to know when it completed its processing, as distinct from the subsequent processing by the originating MTA. In circumstances needing very fine-grained trace information, fields might be created to note all of these "significant" network architecture transitions.

One should note, however, when choosing higher levels of granularity, that the Received header fields present on a message could be counted by MTAs when trying to decide whether or not a message routing loop is in effect. A message with an abundance of these might cause an incorrect determination that the message is in a delivery loop, causing it to be removed from the mail stream. See Section 6.3 of [SMTP] for further discussion.

6. IANA Considerations

6.1. MAIL Parameters Additional-registered-clauses Sub-Registry

This document adds the following entry to the "Additional-registered-clauses" sub-registry of the "MAIL Parameters" registry, created by [SMTP]:

Clause name: state

Description: Indicates entry into a special queue state

Syntax Summary: state <state-name>

Reference: [RFC6729]

6.2. MAIL Parameters Registered-states Sub-Registry

The "MAIL Parameters" registry at IANA has been updated by the creation of the "Registered-states" sub-registry to contain valid state keywords for use with this specification. Updates to this registry are governed by the First Come, First Served rules of [IANA] for new registrations. Changes to the status of existing entries are limited to the original registrant or IESG approval.

Discussion of all registry updates is encouraged via one or more IETF mailing lists that typically cover email-related subjects prior to approval of the change, as a way of documenting the work. The ietf-smtp@ietf.org list is suggested.

Note that only registrations of queue state keywords are permitted. The registry is not to be used for specifying secondary information (i.e., the "value" part of the ABNF in Section 3).

Registrations are to include the following entries:

Name: The name of the state keyword being defined or updated, which conforms to the ABNF shown in Section 3.

Description: A brief description of the keyword's meaning.

Specification: The specification document that defines the queue state being registered, or if no stable reference exists, a more detailed explanation of the queue state than is in the "Description", sufficient to allow interoperability.

Use: One of "current" (the state keyword is in current use), "deprecated" (the state keyword is in use but not recommended for new implementations), or "historic" (the state keyword is no longer in substantial current use).

The initial registration set is as follows:

Name	Description	Specification	Use
auth	Held for message authentication	[RFC6729]	current
content	Held for message content analysis	[RFC6729]	current
convert	Held for message content conversion	[RFC6729]	current
moderation	Held for list moderation	[RFC6729]	current
normal	Message is not being held other than to accommodate typical relaying handling	[RFC6729]	current
other	Held for causes not covered by other registered state keywords	[RFC6729]	current
outbound	Message placed in outbound queue	[RFC6729]	current
quarantine	Held for operator action due to content analysis or local policy	[RFC6729]	current
timed	Held to accommodate a specific requested delivery window	[RFC6729]	current

7. Security Considerations

The use of this trace information can reveal hints as to local policy that was in effect at the time of message handling.

Further discussion about trace field security can be found in Section 7.6 of [SMTP].

8. References

8.1. Normative References

- [IANA] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [MAIL] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.
- [MIME] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.
- [SMTP] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.

8.2. Informative References

- [FUTURERELEASE] White, G. and G. Vaudreuil, "SMTP Submission Service Extension for Future Message Release", RFC 4865, May 2007.

Appendix A. Trace Field Examples

This section includes a sample of the new trace field clause in use.

A.1. Typical Delivery without Obvious Extra Handling

Typical message delivery

```
Received: from newyork.example.com
  (newyork.example.com [192.0.2.250])
  by mail-router.example.net (8.11.6/8.11.6)
  with ESMTTP id i7PK0sH7021929
  for <recipient@example.net>;
  Fri, Feb 15 2002 17:19:22 -0800
Received: from internal.example.com
  (internal.example.com [192.168.0.1])
  by newyork.example.com (8.11.6/8.11.6)
  with ESMTTP id i9MKZCRd064134
  for <recipient@example.net>;
  Fri, Feb 15 2002 17:19:08 -0800
```

Example 1: Typical message delivery with no appreciable extra handling; only Received header fields shown

A.2. Delivery with Moderation

Message delivery after moderation

```
Received: from newyork.example.com
  (newyork.example.com [192.0.2.250])
  by mail-router.example.net (8.11.6/8.11.6)
  with ESMTTP id i7PK0sH7021929
  for <recipient@example.net>;
  Fri, Feb 15 2002 18:33:29 -0800
Received: from internal.example.com
  (internal.example.com [192.168.0.1])
  by newyork.example.com (8.11.6/8.11.6)
  with ESMTTP id i9MKZCRd064134
  for <secret-list@example.com>
  state moderation (sender not subscribed);
  Fri, Feb 15 2002 17:19:08 -0800
```

Example 2: Message held for moderation; only Received header fields shown

The message passed from internal.example.com to newyork.example.com intended for a mailing list hosted at the latter. For list administrative reasons, the message is held there for moderation. It

is finally released over an hour later and passed to the next host. A comment after the state expression indicates the actual cause for the administrative hold.

Appendix B. Acknowledgements

The authors wish to acknowledge the following for their reviews and constructive criticisms of this proposal: Tony Finch, Ned Freed, Carl S. Gutenkunst, John Levine, Bill McQuillan, S. Moonesamy, Alexey Melnikov, Robert A. Rosenberg, Hector Santos, Rolf Sonneveld, and Mykyta Yevstifeyev.

Authors' Addresses

D. Crocker
Brandenburg InternetWorking
675 Spruce Dr.
Sunnyvale 94086
USA

Phone: +1.408.246.8253
EMail: dcrocker@bbiw.net
URI: <http://bbiw.net>

Murray S. Kucherawy
Cloudmark, Inc.
128 King St., 2nd Floor
San Francisco, CA 94107
USA

EMail: superuser@gmail.com

