

Network Working Group
Request for Comments: 3458
Category: Standards Track

E. Burger
SnowShore Networks
E. Candell
Comverse
C. Eliot
Microsoft Corporation
G. Klyne
Nine by Nine
January 2003

Message Context for Internet Mail

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

This memo describes a new RFC 2822 message header, "Message-Context". This header provides information about the context and presentation characteristics of a message.

A receiving user agent (UA) may use this information as a hint to optimally present the message.

Table of Contents

1. Introduction.....	2
2. Conventions used in this document.....	3
3. Motivation.....	3
4. Functional Requirements.....	5
5. Determining the Message Context.....	6
6. Message-Context Reference Field.....	7
6.1. Message-Context Syntax.....	7
6.2. message-context-class Syntax.....	7
6.2.1. voice-message.....	8
6.2.2. fax-message.....	8
6.2.3. pager-message.....	8
6.2.4. multimedia-message.....	8
6.2.5. text-message.....	8
6.2.6. none.....	8
7. Security Considerations.....	9
8. IANA Considerations.....	9
8.1. Message Content Type Registrations.....	9
8.2. Registration Template.....	10
8.3. Message-Context Registration.....	11
9. APPENDIX: Some messaging scenarios.....	12
9.1. Internet e-mail.....	12
9.2. Pager service.....	12
9.3. Facsimile.....	13
9.4. Voice mail.....	14
9.5. Multimedia message.....	14
10. References.....	15
10.1 Normative References.....	15
10.2 Informative References.....	15
11. Acknowledgments.....	15
12. Authors' Addresses.....	16
13. Full Copyright Statement.....	17

1. Introduction

This document describes a mechanism to allow senders of an Internet mail message to convey the message's contextual information. Taking account of this information, the receiving user agent (UA) can make decisions that improve message presentation for the user in the context the sender and receiver expects.

In this document, the "message context" conveys information about the way the user expects to interact with the message. For example, a message may be e-mail, voice mail, fax mail, etc. A smart UA may have specialized behavior based on the context of the message.

This document specifies a RFC 2822 header called "Message-Context".

The mechanism is in some ways similar to the use of the Content-Disposition MIME entity described in [6]. Content-Disposition gives clues to the receiving User Agent (UA) for how to display a given body part. Message-Context can give clues to the receiving UA for the presentation of the message. This allows the receiving UA to present the message to the recipient, in a meaningful and helpful way.

Typical uses for this mechanism include:

- o Selecting a special viewer for a given message.
- o Selecting an icon indicating the kind of message in a displayed list of messages.
- o Arranging messages in an inbox display.
- o Filtering messages the UA presents when the user has limited access.

2. Conventions used in this document

This document refers generically to the sender of a message in the masculine (he/him/his) and the recipient of the message in the feminine (she/her/hers). This convention is purely for convenience and makes no assumption about the gender of a message sender or recipient.

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 BCP 14, RFC 2119 [2].

FORMATTING NOTE: Notes, such as this one, provide additional nonessential information that the reader may skip without missing anything essential. The primary purpose of these non-essential notes is to convey information about the rationale of this document, or to place this document in the proper historical or evolutionary context. Readers whose sole purpose is to construct a conformant implementation may skip such information. However, it may be of use to those who wish to understand why we made certain design choices.

3. Motivation

Multimedia messaging systems receive messages that a UA may present in variety of ways. For example, traditional e-mail uses simple text messages that the recipient displays and edits. One UA may automatically print Fax images. Another UA may play voice messages through a telephone handset. Likewise, a receiving desktop computer

may process or present documents transferred over e-mail using a local application. Emerging and future developments may deliver other forms of information that have their own characteristics for user presentation, such as video messages and pager messages.

An often-requested characteristic for multimedia messaging systems is to collect received messages in a "universal inbox", and to offer them to the user as a combined list.

In the context of "unified messaging", different message contexts may have different implied semantics. For example, some users may perceive voicemail to have an implicit assumption of urgency. Thus they may wish to gather them together and process them before other messages. This results in the end-user receiving agent needing to be able to identify voicemail and distinguish it from other messages.

The uses of this kind of presentation characteristic for each message are multi-fold:

- o Display an indication to the user (e.g., by a suitably evocative icon along with other summary fields),
- o Auto-forward a given message type into another messaging environment (e.g., a page to a mobile short message service),
- o Prioritize and group messages in an inbox display list,
- o Suggest appropriate default handling for presentation,
- o Suggest appropriate default handling for reply, forward, etc.

A problem faced by multimedia messaging systems is that it is not always easy to decide the context of a received message. For example, consider the following scenarios.

- o A message that contains audio and image data: Is this a fax message that happens to have some voice commentary? Is it a voice message that is accompanied by some supplementary diagrams? Is it a fully multimedia message, in which all parts are expected to carry equal significance?
- o A message containing text and audio data: Is this e-mail with an MP3 music attachment? Is it a voice message that happens to have been generated with an initial text header for the benefit of non-voice-enabled e-mail receivers?

The message context does relate to the message media content. However, it is not the same thing. As shown above, the media type used in a message is not sufficient to indicate the message context. One cannot determine a priori which media types to use in alternative (gateway) messages. Also, what if the user cares about distinguishing traditional e-mail text from SMS messages? They are both the same media type, text, but they have different user contexts.

4. Functional Requirements

The goals stated above lead to the following functional requirements.

For receivers:

- o Identify a message as belonging to a message class.
- o Incorrect or invalid message classification must not result in failure to transfer or inability to present a message.

For senders:

- o Specify message classes by the originating user's choice of authoring tool or simple user interaction.

For both:

- o Specify a well-defined set of message classes to make interoperability between mail user agents (UAs) possible.
- o Message classification information has to be interpretable in reasonable fashion by many different user agent systems.
- o The mechanism should be extensible to allow for the introduction of new kinds of messages.

NOTE: We specifically do not specify user agent behavior when the user agent forwards a message. Clearly, the user agent, being message-context-aware, should provide a meaningful message-context. It is obvious what to do for the easy cases. Messages that the user simply forwards will most likely keep the context unchanged. However, it is beyond the scope of this document to specify the user agent behavior for any other scenario.

5. Determining the Message Context

One method of indicating the interpretation context of a message is to examine the media types in the message. However, this requires the UA to scan the entire message before it can make this determination. This approach is particularly burdensome for the multi-media mail situation, as voice and especially video mail objects are quite large.

We considered indicating the message context by registering a multipart/* MIME subtype (Content-Type). For example, the VPIM Work Group has registered multipart/voice-message to indicate that a message is primarily voice mail [7]. However, multipart/voice-message is identical in syntax to multipart/mixed. The only difference is that VPIM mail transfer agents and user agents recognize that they can perform special handling of the message based on it being a voice mail message. Moreover, Content-Type refers to a given MIME body part, not to the message as a whole.

We wish to avoid scanning the entire message. In addition, we wish to avoid having to create multiple aliases for multipart/mixed every time someone identifies a new primary content type. Multiple aliases for multipart/mixed are not desirable as they remove the possibility for specifying a message as multipart/alternate, multipart/parallel, or multipart/encrypted, for example.

Since the message context is an attribute of the entire message, it is logical to define a new top-level (RFC 2822 [3]) message attribute. To this end, this document introduces the message attribute "Message-Context".

Message-Context only serves to identify the message context. It does not provide any indication of content that the UA must be capable of delivering. It does not imply any message disposition or delivery notification. There is a related effort to define Critical Content of Internet Mail [8] that one might use to perform these tasks.

Message-Context is only an indicator. We do not intend for it to convey information that is critical for presentation of the message. One can conceive of goofy situations, such as a message marked "voice-message" but without an audio body part. In this case, the fact that the contents of a message don't match its context does not mean the receiving system should generate an error report or fail to deliver or process the message.

6. Message-Context Reference Field

The Message-Context reference field is a top-level header inserted by the sending UA to indicate the context of the message.

A receiving user agent MUST NOT depend on the indicated message-context value in a way that prevents proper presentation of the message. If the value is incorrect or does not match the message content, the receiving user agent MUST still be capable of displaying the message content at least as meaningfully as it would if no Message-Context value were present.

One can envision situations where a well-formed message ends up not including a media type one would expect from the message-context. For example, consider a voice messaging system that records a voice message and also performs speech-to-text processing on the message. The message then passes through a content gateway, such as a firewall, that removes non-critical body parts over a certain length. The receiving user agent will receive a message in the voice-message context that has only a text part and no audio. Even though the message does not have audio, it is still in the voice message context.

Said differently, the receiving UA can use the message-context to determine whether, when, and possibly where to display a message. However, the message-context should not affect the actual rendering or presentation. For example, if the message is in the voice-message context, then don't try to send it to a fax terminal. Conversely, consider the case of a message in the voice-message context that gets delivered to a multimedia voice terminal with a printer. However, this message only has fax content. In this situation, the "voice-message" context should not stop the terminal from properly rendering the message.

6.1. Message-Context Syntax

The syntax of the Message-Context field, described using the ABNF [4] is as follows. Note that the Message-Context header field name and message-context-class values are not case sensitive.

```
"Message-Context" ":" message-context-class CRLF
```

6.2. message-context-class Syntax

The message-context-class indicates the context of the message. This is an IANA registered value. Current values for message-context-class are as follows.

```
message-context-class = ( "voice-message"  
                          / "fax-message"  
                          / "pager-message"  
                          / "multimedia-message"  
                          / "text-message"  
                          / "none"  
                          )
```

Note: The values for Message-Context MUST be IANA registered values following the directions in the IANA Considerations section below.

6.2.1. voice-message

The voice-message class states the message is a voice mail message.

6.2.2. fax-message

The fax-message class states the message is a facsimile mail message.

6.2.3. pager-message

The pager-message class states the message is a page, such as a text or numeric pager message or a traditional short text message service (SMS) message.

6.2.4. multimedia-message

The multimedia-message class states the message is an aggregate multimedia message, such as a message specified by [9]. This helps identify a message in a multimedia context. For example, a MIME multipart/related [10] data part and resource part looks the same as a multimedia MHTML multipart/related. However, the semantics are quite different.

6.2.5. text-message

The text-message class states the message is a traditional internet mail message. Such a message consists of text, possibly richly formatted, with or without attachments.

6.2.6. none

The none class states there is no context information for this message.

If a message has no Message-Context reference field, a receiving user agent MUST treat it the same as it would if the message has a "none" value.

7. Security Considerations

The intention for this header is to be an indicator of message context only. One can imagine someone creating an "Application" Message-Context. A poorly designed user agent could blindly execute a mailed program based on the Message-Context. Don't do that!

One can envision a denial of service attack by bombing a receiver with a message that has a Message-Context that doesn't fit the profile of the actual body parts. This is why the receiver considers the Message-Context to be a hint only.

8. IANA Considerations

Section 8.3 is a registration for a new top-level RFC 2822 [3] message header, "Message-Context".

This document creates an extensible set of context types. To promote interoperability and coherent interpretations of different types, a central repository has been established for well-known context types.

The IANA has created a repository for context types called "Internet Message Context Types". Following the policies outlined in [5], this repository is "Specification Required" by RFC. Section 8.1 describes the initial values for this registry.

To create a new message context type, you MUST publish an RFC to document the type. In the RFC, include a copy of the registration template found in Section 8.2 of this document. Put the template in your IANA Considerations section, filling-in the appropriate fields. You MUST describe any interoperability and security issues in your document.

8.1. Message Content Type Registrations

Internet Message Content Types

=====

Value	Description	Reference
-----	-----	-----
voice-message	Indicates a message whose primary content is a voice mail message. The primary content is audio data. The context is usually a message recorded from a voice telephone call.	This RFC

fax-message	Indicates a message whose primary content is a fax mail message. The primary content is image data. The context is usually a message recorded from a facsimile telephone call.	This RFC
pager-message	Indicates a message whose primary content is a page. The primary content is text data. The context is an urgent message usually of a limited length.	This RFC
multimedia-message	Indicates a message whose primary content is a multimedia message. The primary content is multimedia, most likely MHTML. The context is often spam or newsletters.	This RFC
text-message	Indicates a classic, text-based, Internet message.	This RFC
None	Indicates an unknown message context.	This RFC

8.2. Registration Template

In the following template, a pipe symbol, "|", precedes instructions or other helpful material. Be sure to replace "<classname>" with the class name you are defining.

Message-Context class name:
<classname>

Summary of the message class:

| Include a short (no longer than 4 lines) description or summary

| Examples:

| "Palmtop devices have a 320x160 pixel display, so we can..."

| "Color fax is so different than black & white that..."

Person & email address to contact for further information:

| Name & e-mail

8.3. Message-Context Registration

To: iana@iana.org
Subject: Registration of New RFC 2822 Header

RFC 2822 Header Name:
Message-Context

Allowable values for this parameter:
Please create a new registry for Primary Context Class
registrations. See section 8.1 of this document for the initial
values.

RFC 2822 Section 3.6 Repeat Value:

Field	Min Number	Max Number	Notes
Message-Context	0	1	

Person & email address to contact for further information:
Eric Burger
e.burger@ieee.org

9. APPENDIX: Some messaging scenarios

This section is not a normative part of this document. We include it here as a historical perspective on the issue of multimedia message types.

These scenarios are neither comprehensive nor fixed. For example, e-mails being typically text-based do not mean that they cannot convey a voice-message. This very mutability serves to underline the desirability of providing some explicit message context hint.

9.1. Internet e-mail

Internet e-mail carries textual information. Sometimes it conveys computer application data of arbitrary size.

Typically, one uses e-mail for non-urgent messages, which the recipient will retrieve and process at a time convenient to her.

The normal device for receiving and processing e-mail messages is some kind of personal computer. Modern personal computers usually come with a reasonably large display and an alphanumeric keyboard. Audio, video, and printing capabilities are not necessarily available.

One can use E-mail for communication between two parties (one-to-one), a small number of known parties (one-to-few) or, via an e-mail distribution list, between larger numbers of unknown parties (one-to-many).

One of the endearing characteristics of e-mail is the way that it allows the recipient to forward all or part of the message to another party, with or without additional comments. It is quite common for an e-mail to contain snippets of content from several previous messages. Similar features apply when replying to e-mail.

9.2. Pager service

One uses a pager message to convey notifications and alerts. For the most part, these notifications are textual information of limited size. The typical limit is 160 characters. People use pages for relatively urgent messages, which the sender wishes the receiver to see and possibly respond to within a short time period. Pager messages are often used as a way of alerting users to something needing their attention. For example, a system can use a page to notify a subscriber there is a voicemail message requiring her attention.

Example devices for sending and receiving a pager message are a mobile telephone with a small character display or a text pager. Personal computers and personal digital assistants (PDAs) can also participate in pager messaging.

Currently, the most common use of pager messages are between just two parties (one-to-one).

One delivery method for pager messages is the short text messaging service (SMS). SMS is a facility that has evolved for use with mobile telephones, and has an associated per-message transmission charge. Note that the focus here is on the notification aspect of SMS. From the beginning, SMS was envisioned to be more than a simple pager service. Operators can use SMS to provision the phone, for example. From the subscriber point of view, SMS has evolved considerably from its origins as a pure pager replacement service. For example, with mobile originate service, people can have two-way text chat sessions using SMS and a mobile phone. In addition, there are SMS-enabled handsets that can display pictures. However, for the purposes of this document, there is still a need to capture the essence of a "highly urgent, short-text, notification or alert" service.

Users often send pager messages in isolation, rather than as part of a longer exchange. One use for them is as a prompt or invitation to communicate by some more convenient and content-rich method, such as a telephone call.

9.3. Facsimile

People use facsimile to convey image information of moderate size, typically a small number of pages. Sometimes people use facsimile for larger documents.

Facsimile is a facility that usually uses circuit-switched telephone circuits, with connection-time charges. Message transfer takes place in real-time. Thus, people often use facsimile for urgent communication.

The normal device for sending and receiving a facsimile is a self-contained scanning and printing device connected to a telephone line or a desktop computer.

Most facsimiles are between just two parties (one-to-one). However, a significant portion of facsimile service is broadcast between multiple parties (one-to-many).

Most facsimile exchanges are in isolation, rather than as part of a longer exchange. Facsimile data is typically not suitable for further processing by computer.

9.4. Voice mail

People use voice mail to convey audio information, almost exclusively human speech.

Voice mail is a facility that usually uses circuit-switched telephone circuits, with modest connection-time charges, often used for moderately urgent messages. A common use for them is as a prompt or invitation to communicate by some more convenient method, such as a telephone call. In most, but not all cases, the sender of a voice message does not want to send a message at all. Rather, they wished to engage in a real-time conversation.

The normal device for sending and receiving a voice mail is a telephone handset.

Voice messages are usually sent between just two parties (one-to-one).

Voice mail data is not generally suitable for further processing by computer.

9.5. Multimedia message

We define a multimedia message as a message containing more than one basic media type (text, image, audio, video, model, application).

The following are some characteristics of a multimedia message.

In some cases, a multimedia message is just e-mail with an attachment that a multimedia display application presents. For example, I can send you an MP3 of something I recorded in my garage today.

In other cases, a multimedia message represents a convergence between two or more of the scenarios described above. For example, a voice message with an accompanying diagram or a talking head video message is a multimedia message.

The characteristics will vary somewhat with the intent of the sender. This in turn may affect the user agent or application used to render the message.

10. References

10.1 Normative References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Resnick, P., "Internet Message Format", RFC 2822, April 2001.
- [4] Crocker, D. and P. Overell, Eds., "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, November 1997.
- [5] Alvestrand, H. and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

10.2 Informative References

- [6] Troost, R., Dorner, S. and K. Moore, "Communicating Presentation Information in Internet Messages: The Content-Disposition Header Field", RFC 2183, August 1997.
- [7] Vaudreuil, G. and G. Parsons, "VPIM Voice Message MIME Sub-type Registration", RFC 2423, September 1998.
- [8] Burger, E., "Critical Content of Internet Mail", RFC 3459, January 2003.
- [9] Palme, J., Hopmann, A. and N. Shelness, "MIME Encapsulation of Aggregate Documents, such as HTML (MHTML)", RFC 2557, March 1999.
- [10] Levinson, E., "The MIME Multipart/Related Content-type", RFC 2387, August 1998.

11. Acknowledgments

Many of the ideas here arose originally from a discussion with Jutta Degener.

We'd also like to thank Keith Moore for helping us tighten-up our explanations.

In the last round, we got some rather good advise from Caleb Clausen and Dave Aronson.

Antti Vaha-Sipila pointed out advances in SMS, while Stuart McRae helped distil the essence of the pager service vis a vis SMS.

We offer an extra special thanks to Greg Vaudreuil for pulling RFC 2557 out of his hat.

12. Authors' Addresses

Eric Burger
SnowShore Networks, Inc.
285 Billerica Rd.
Chelmsford, MA 01824-4120
USA

Phone: +1 978 367 8403
EMail: e.burger@ieee.org

Emily Candell
Comverse Network Systems
200 Quannapowitt Pkwy.
Wakefield, MA 01880
USA

Phone: +1 781 213 2324
EMail: emily.candell@comverse.com

Graham Klyne
Nine by Nine
United Kingdom

EMail: GK-IETF@ninebynine.org

Charles Eliot
Microsoft Corporation
One Microsoft Way
Redmond WA 98052
USA

Phone: +1 425 706 9760
EMail: charle@Microsoft.com

13. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.