

Workgroup: MIMI BoF
Internet-Draft: draft-mahy-mimi-content-01
Published: 24 October 2022
Intended Status: Informational
Expires: 27 April 2023
Authors: R. Mahy
Wire

More Instant Messaging Interoperability (MIMI) message content

Abstract

This document describes content semantics common in Instant Messaging (IM) systems and describes an example profile suitable for instant messaging interoperability of messages end-to-end encrypted inside the MLS (Message Layer Security) Protocol. It adapts prior work (CPIM) to work well in the MLS context.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 27 April 2023.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Terminology](#)
- [2. Introduction](#)
- [3. Overview](#)
 - [3.1. Naming schemes](#)
 - [3.2. Detection of suitable media types](#)
 - [3.3. CPIM and MIME headers](#)
- [4. Example](#)
 - [4.1. Original Message](#)
 - [4.2. Reply](#)
 - [4.3. Reaction](#)
 - [4.4. Mentions](#)
 - [4.5. Edit](#)
 - [4.6. Delete](#)
 - [4.7. Unlike](#)
 - [4.8. Expiring](#)
 - [4.9. Knock](#)
 - [4.10. Read Receipt](#)
 - [4.11. Attachments](#)
 - [4.12. Conferencing](#)
- [5. IMMI CPIM profile](#)
 - [5.1. CPIM headers](#)
 - [5.2. Definition of message/immi-disposition-notification](#)
 - [5.3. Required and Recommended MIME types](#)
- [6. IANA Considerations](#)
 - [6.1. MIME subtype registration of message/immi-disposition-notification](#)
- [7. Security Considerations](#)
- [8. Normative References](#)
- [9. Informative References](#)
- [Appendix A. TLS Presentation Language multipart container format Author's Address](#)

1. Terminology

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 RFC 2119 [[RFC2219](#)].

The terms MLS client, MLS group, and KeyPackage have the same meanings as in the MLS protocol [[I-D.ietf-mls-protocol](#)].

2. Introduction

MLS [[I-D.ietf-mls-protocol](#)] is a group key establishment protocol motivated by the desire for group chat with efficient end-to-end encryption. While one of the motivations of MLS is interoperable standards-based secure messaging, the MLS protocol does not define or

prescribe any format for the encrypted "application messages" encoded by MLS. The development of MLS was strongly motivated by the needs of a number of Instant Messaging (IM) systems, which encrypt messages end-to-end using variations of the Double Ratchet protocol [[DoubleRatchet](#)].

End-to-end encrypted instant messaging was also a motivator for the Common Protocol for Instant Messaging (CPIM) [[RFC3862](#)], however the model used at the time assumed standalone encryption of each message using a protocol such as S/MIME [[RFC8551](#)] or PGP [[RFC3156](#)] to interoperate between IM protocols such as SIP [[RFC3261](#)] and XMPP [[RFC6120](#)]. For a variety of practical reasons, interoperable end-to-end encryption between IM systems was never deployed commercially.

There are now several instant messaging vendors implementing MLS. In order to enable interoperable messaging conveyed "inside" MLS application messages, some additional specification and some minor changes are required. Also, the expectation of what constitutes basic features common across multiple IM systems has grown. It would be beneficial to provide an interoperable format for these additional features as well. Most of these features could be implemented using a profile which describes how to use already-defined URIs, message headers, and MIME types.

This document explores issues and example solutions consistent with the More Instant Messaging Interoperability (MIMI) problem outline [[I-D.mahy-mimi-problem-outline](#)].

This proposal assumes that MLS clients can advertise media types they support and that MLS clients can determine what media types are required to join a specific MLS group. Specifically, [[I-D.mahy-mls-content-adv](#)] defines two MLS extensions which meet this requirement. It would allow implementations to define groups with different MIME type requirements and it would allow MLS clients to send extended or proprietary messages that would be interpreted by some members of the group while assuring that an interoperable end-to-end encrypted baseline is available to all members, even when the group spans multiple systems or vendors.

Below is a list of some features commonly found in IM group chat systems:

- *plain text and rich text messaging
- *delivery notifications
- *read receipts
- *replies
- *reactions
- *edit or delete previously sent messages
- *expiring messages

- *knock / ping
- *shared files/audio/videos
- *calling / conferencing

3. Overview

3.1. Naming schemes

IM systems have a number of types of identifiers. These are described in detail in [[I-D.mahy-mimi-identity](#)]. A few of these used in this document are:

- *handle identifier (external, friendly representation). This is the type of identifier in the From header in the examples.
- *client/device identifier (internal representation). This is the type of identifier in the implied Sender header in the examples.
- *group or conversation or channel name (either internal or external representation). . This is the type of identifier in the implied To header in the examples.

This proposal relies on URIs for naming and identifiers. All the example use the im: URI scheme (defined in [[RFC3862](#)]), but any instant messaging scheme could be used.

3.2. Detection of suitable media types

As most IM systems are proprietary, standalone systems, it is useful to allow clients to send and receive proprietary formats among themselves. Using the multipart/alternative MIME wrapper (or the container syntax in Appendix A of this document), clients can send a message using the basic functionality described in this document AND a proprietary format for same-vendor clients simultaneously over the same group with end-to-end encryption.

Example sending this profile and proprietary messaging protocol simultaneously.

Content-type: multipart/alternative; boundary=XcrSXMwuRwk9

--XcrSXMwuRwk9

Content-type: message/cpim

From: <im:alice-smith@example.com>

DateTime: 2022-02-08T22:13:45-00:00

Message-ID: <28fd19857ad7@example.com>

Content-Type: text/plain; charset=utf-8

Test Message

--XcrSXMwuRwk9

Content-type: application/vnd.examplevendor-fancy-im-message

<content of example vendor's fancy proprietary format>

--XcrSXMwuRwk9

[[I-D.mahy-mls-content-adv](#)] contains the actual MLS extensions useful for advertising the relevant media types. The profile in this document requires support for receiving:

- *message/cpim
- *text/plain
- *text/markdown
- *multipart/mixed
- *multipart/alternative

All other mime types (including some recommended in this profile) are optional.

3.3. CPIM and MIME headers

Note that while the **syntax** used in this document is based on CPIM, the semantics of these messages could be translated into any number of formats, for example JSON, XML, CBOR, etc.

We assume that an MLS group is already established and that either out-of-band or using the MLS protocol or MLS extensions that the following is known to every member of the group:

- *The membership of the group (via MLS).
- *The identity of any MLS client which sends an application message (via MLS).
- *The MLS group ID (via MLS)
- *The human readable name(s) of the MLS group, if any (out-of-band or extension).
- *Which media types are mandatory to implement (proposed extension).
- *For each member, the media types each supports (proposed extension).

For all messages the message header equivalent of To (the MLS group) and Sender fields (MLS sender) is already known and is therefore redundant. Every message contains a message/cpim header which includes the From, DateTime, and Message-ID fields. The From field contains the external, user-friendly representation of the Sender.

Messages sent to an MLS group are delivered to every member of the group active during the epoch in which the message was sent.

It is also mandatory to understand are the following MIME headers:

- *Content-Type
- *Content-Disposition
- *Content-Length

4. Example

4.1. Original Message

In this example, Alice Smith sends a rich-text (Markdown) [[RFC7763](#)] message to the Engineering Team MLS group. The following values are implied as if headers were present:

- *Implied Sender header from MLS sender: im:3b52249d-68f9-45ce-8bf5-c799f3cad7ec/0003@example.com
- *Implied To header from MLS group: "Engineering Team"
im:#engineering_team@example.com

Content-type: message/cpim

From: <im:alice-smith@example.com>
DateTime: 2022-02-08T22:13:45-00:00
Message-ID: <28fd19857ad7@example.com>

Content-Type: text/markdown;charset=utf-8

Hi everyone, we just shipped release 2.0. __Good work__!

4.2. Reply

A reply message looks similar, but contains an In-Reply-To CPIM header with the ID of the original message. The implied To header is the same all example messages in this section. The implied Sender header is always the MLS sender, and will not be shown in subsequent example messages.

Content-type: message/cpim

From: <im:bob-jones@example.com>
DateTime: 2022-02-08T22:13:57-00:00
Message-ID: <e701beee59f9@example.com>
In-Reply-To: <28fd19857ad7@example.com>

Content-Type: text/markdown;charset=utf-8

Right on! _Congratulations_ 'all!

4.3. Reaction

A reaction, uses the reaction Content-Disposition token defined in [\[RFC9078\]](#). This Content-Disposition token indicates that the intended disposition of the contents of the message is a reaction.

The content in the sample message is a single Unicode heart character (U+2665). Discovering the range of characters each implementation could render as a reaction can occur out-of-band and is not within the scope of this proposal. However, an implementation which receives a reaction character string it does not recognize could render the reaction as a reply, possibly prefixing with a localized string such as "Reaction: ". Note that a reaction could theoretically even be another media type (ex: image, audio, or video), although not currently implemented in major instant messaging systems.

Content-type: message/cpim

From: <im:cathy-washington@example.com>
DateTime: 2022-02-08T22:13:57-00:00
Message-ID: <1a771ca1d84f@example.com>
In-Reply-To: <28fd19857ad7@example.com>

Content-Type: text/plain;charset=utf-8
Content-Disposition: reaction

♥

4.4. Mentions

In instant messaging systems and social media, a mention allows special formatting and behavior when a name, handle, or tag associated with a known group is encountered, often when prefixed with a commercial-at "@" character for mentions of users or a hash "#" character for groups or tags. A message which contains a mention may trigger distinct notifications on the IM client.

We can convey a mention by linking the user, handle, or tag URI in Markdown or HTML rich content. For example, a mention using Markdown is indicated below.

Content-type: message/cpim

From: <im:cathy-washington@example.com>
DateTime: 2022-02-08T22:14:03-00:00
Message-ID: <4dcab7711a77@example.com>

Content-Type: text/markdown;charset=utf-8

Kudos to [[@Alice Smith](mailto:alice-smith@example.com)](im:alice-smith@example.com)
for making the release happen!

The same mention using HTML [[W3C.CR-html52-20170808](#)] is indicated below.

Content-type: message/cpim

From: <im:cathy-washington@example.com>
DateTime: 2022-02-08T22:14:03-00:00
Message-ID: <4dcab7711a77@example.com>

Content-Type: text/html;charset=utf-8

<p>Kudos to @Alice
Smith for making the release happen!</p>

4.5. Edit

Unlike with email messages, it is common in IM systems to allow the sender of a message to edit or delete the message after the fact. Typically the message is replaced in the user interface of the receivers (even after the original message is read) but shows a visual indication that it has been edited.

We reuse the Supersedes header from MIXER [[RFC2156](#)], because the semantics are correct: the message included in the body is a replacement for the message with the superseded message ID.

Here Bob Jones corrects a typo in his original message:

Content-type: message/cpim

From: <im:bob-jones@example.com>
DateTime: 2022-02-08T22:13:57-00:00
Message-ID: <89d3472622a4@example.com>
Supersedes: <e701beee59f9@example.com>

Content-Type: text/markdown;charset=utf-8

Right on! _Congratulations_ y'all!

4.6. Delete

In IM systems, a delete means that the author of a specific message has retracted the message, regardless if other users have read the message or not. Typically a placeholder remains in the user interface showing that a message was deleted. Replies which reference a deleted message typically hide the quoted portion and reflect that the original message was deleted.

If Bob deleted his message instead of modifying it, we would represent it using the Supersedes header with an empty body, as shown below.

Content-type: message/cpim

From: <im:bob-jones@example.com>
DateTime: 2022-02-08T22:13:57-00:00
Message-ID: <89d3472622a4@example.com>
Supersedes: <e701beee59f9@example.com>

Content-Length: 0

4.7. Unlike

In most IM systems, not only is it possible to react to a message ("Like"), but it is possible to remove a previous reaction ("Unlike"). This can be accomplished by deleting the message which creates the original reaction

If Cathy removes her reaction, we would represent the removal using a Supersedes header with an empty body, referring to the message which created the reaction, as shown below.

Content-type: message/cpim

From: <im:cathy-washington@example.com>
DateTime: 2022-02-08T22:14:14-00:00
Message-ID: <d052cace46f8@example.com>
Supercedes: <1a771ca1d84f@example.com>

Content-Length: 0

4.8. Expiring

Expiring messages are designed to be deleted automatically by the receiving client at a certain time whether they have been read or not. As with manually deleted messages, there is no guarantee that a uncooperative client or a determined user will not save the content of the message, however most clients respect the convention.

MIXER defines an Expires header which is also used sent simply by including an Expires header in the CPIM message body.

To avoid using two different date header syntaxes, we define an ExpiresDateTime header, which uses the same date/time format as CPIM's DateTime header. The semantics of the header are that the message is automatically deleted by the receiving clients at the indicated time without user interaction or network connectivity necessary.

Content-type: message/cpim

From: <im:alice-smith@example.com>
DateTime: 2022-02-08T22:49:03-00:00
Message-ID: <5c95a4dfddab@example.com>
ExpiresDateTime: 2022-02-08T22:59:03-00:00

Content-Type: text/markdown;charset=utf-8

___*VPN GOING DOWN*___

I'm rebooting the VPN in ten minutes unless anyone objects.

4.9. Knock

A knock or ping is message sent to get the attention of a user or a group of users. It might be sent when a user has not responded to direct messages or mentions, or in a group when something requires the attention of everyone quickly (ex: a serious unusual situation like a major system outage).

We represent a knock as a text/plain body containing a single CRLF with the alert Content-Disposition token (defined in [[RFC3261](#)]).

Content-type: message/cpim

From: <im:alice-smith@example.com>
DateTime: 2022-02-08T22:13:45-00:00
Message-ID: <c1a3375bfe3f@example.com>

Content-Type: text/plain
Content-Disposition: alert

4.10. Read Receipt

In instant messaging systems, read receipts typically generate a distinct indicator for each message. In some systems, the number of users in a group who have read the message is subtly displayed and the list of users who read the message is available on further inspection.

Of course, Internet mail has support for read receipts as well, but the existing message disposition notification mechanism defined for email in [[RFC8098](#)] is unfortunately inappropriate in this context.

- *notifications can be sent by intermediaries
- *only one notification can be sent about a single message per recipient
- *a human-readable version of the notification is expected
- *each notification can refer to only one message
- *it is extremely verbose

The proposed format below, message/immi-disposition-notification is sent by one member of an MLS group to the entire group and can refer to multiple messages. There is one IMMI-Disposition line per message, with the disposition of the original message in a parameter. As the disposition at the recipient changes, the disposition can be updated in a subsequent notification.

Content-type: message/cpim

From: <im:bob-jones@example.com>
DateTime: 2022-02-09T07:57:13-00:00
Message-ID: <7e924c2e6ee5@example.com>

Content-Disposition: notification
Content-type: message/immi-disposition-notification

IMMI-Disposition: <4dcab7711a77@example.com>;dispo=read
IMMI-Disposition: <285f75c46430@example.com>;dispo=read
IMMI-Disposition: <c5e0cd6140e6@example.com>;dispo=read
IMMI-Disposition: <5c95a4dfddab@example.com>;dispo=expired

4.11. Attachments

The message/external-body MIME Type is a convenient way to present a URL to download an attachment which should not be rendered inline.

```
Content-Type: message/external-body; access-type="URL";
URL="https://example.com/storage/bigfile.m4v";
size=708234961
```

4.12. Conferencing

Joining a conference via URL is also possible. The link could be rendered to the user, requiring a click. Alternatively another Content-Disposition could be specified to more automatic actions. However further calling and conferencing functionality is out-of-scope of this document.

```
Content-Type: message/external-body; access-type="URL";
URL="https://example.com/join/12345"
```

5. IMMI CPIM profile

We define a profile of CPIM for instant messaging within MLS. The grammar uses Augmented Backus-Naur Form (BNF) [[RFC5234](#)].

5.1. CPIM headers

The following CPIM headers are required:

- *From: the identity of message sender. for example im:alice@example.com this identity could be pseudonymous or anonymous if the group policy allows.
- *DateTime: the date and time in a reasonable format, as specified in CPIM.
- *Message-ID: a message ID which is unique across domains.
- *Content-type: As is from CPIM.
- *In-Reply-To: Refers to the previous Message-ID. Same semantics as in [[RFC5322](#)].
- *Supersedes: Refers to the previous Message-ID. Similar semantics to header of the same name in MIXER. Content-Disposition: The intended handling of the message. The two required dispositions are render and reaction.
- *Content-Length:

For clarity the grammar for the headers not already included in CPIM are formulated below.

```

msg-id-header-line = msg-id-header ":" SP msg-id CRLF
msg-id-header = "Message-ID" ; case-sensitive

in-reply-to-header-line = in-reply-to-header ":" SP msg-id CRLF
in-reply-to-header = "In-Reply-To" ; case-sensitive

supersedes-header-line = supersedes-header ":" SP msg-id CRLF
supersedes-header = "Supersedes" ; case-sensitive

msg-id = "<" id-left "@" id-right ">"

id-left = dot-atom-text
id-right = dot-atom-text / no-fold-literal

dot-atom-text = 1*atext *("." 1*atext)

atext = ALPHA / DIGIT / atom-symbol

atom-symbol = "!" / "#" / "$" / "%" / "&" / "'" / "*" / "+" / "-" /
              "/" / "=" / "?" / "^" / "_" / "`" / "{" / "|" / "}" / "~"

no-fold-literal = "[" *dtext "]"

dtext = %d33-90 / %d94-126 ; Printable US-ASCII
        ; excluding "[", "]", and "\"

```

5.2. Definition of message/immi-disposition-notification

The grammar below defines the syntax.

```

immi-disposition-notification-body = 1*immi-header-line

immi-header-line = immi-header ":" SP msg-id ";" status CRLF

immi-header = "IMMI-Disposition" ; case-sensitive

status = "dispo" "=" status-value

status-value = "read" /
              "error" /
              "delivered" /
              "expired" /
              "deleted" /
              "hidden"

```

5.3. Required and Recommended MIME types

The following MIME types are REQUIRED:

```
*message/cpim
```

- *multipart/alternative
- *multipart/mixed
- *multipart/parallel
- *text/plain
- *text/markdown

The following MIME types are RECOMMENDED:

- *text/html
- *message/external-body
- *message/immi-disposition-notification
- *image/jpeg
- *image/png

6. IANA Considerations

6.1. MIME subtype registration of message/immi-disposition-notification

This document proposes registration of a MIME subtype with IANA.

TBC

7. Security Considerations

TBC

8. Normative References

[I-D.mahy-mimi-problem-outline]

Mahy, R., "More Instant Messaging Interoperability (MIMI) problem outline", Work in Progress, Internet-Draft, draft-mahy-mimi-problem-outline-01, 24 October 2022, <<https://datatracker.ietf.org/api/v1/doc/document/draft-mahy-mimi-problem-outline/>>.

[I-D.mahy-mls-content-adv]

Mahy, R., "Content Type Advertisement for Message Layer Security (MLS)", Work in Progress, Internet-Draft, draft-mahy-mls-content-adv-00, 23 October 2022, <<https://www.ietf.org/archive/id/draft-mahy-mls-content-adv-00.txt>>.

[RFC2156] Kille, S., "MIXER (Mime Internet X.400 Enhanced Relay): Mapping between X.400 and RFC 822/MIME", RFC 2156, DOI 10.17487/RFC2156, January 1998, <<https://www.rfc-editor.org/info/rfc2156>>.

[RFC2219] Hamilton, M. and R. Wright, "Use of DNS Aliases for Network Services", BCP 17, RFC 2219, DOI 10.17487/RFC2219, October 1997, <<https://www.rfc-editor.org/info/rfc2219>>.

- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<https://www.rfc-editor.org/info/rfc3261>>.
- [RFC3862] Klyne, G. and D. Atkins, "Common Presence and Instant Messaging (CPIM): Message Format", RFC 3862, DOI 10.17487/RFC3862, August 2004, <<https://www.rfc-editor.org/info/rfc3862>>.
- [RFC5234] 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>>.
- [RFC7763] Leonard, S., "The text/markdown Media Type", RFC 7763, DOI 10.17487/RFC7763, March 2016, <<https://www.rfc-editor.org/info/rfc7763>>.

9. Informative References

- [DoubleRatchet] Perrin, T. and M. Marlinspike, "The Double Ratchet Algorithm", 20 November 2016, <<https://signal.org/docs/specifications/doubleratchet/>>.
- [I-D.ietf-mls-protocol] Barnes, R., Beurdouche, B., Robert, R., Millican, J., Omara, E., and K. Cohn-Gordon, "The Messaging Layer Security (MLS) Protocol", Work in Progress, Internet-Draft, draft-ietf-mls-protocol-16, 11 July 2022, <<https://www.ietf.org/archive/id/draft-ietf-mls-protocol-16.txt>>.
- [I-D.mahy-mimi-identity] Mahy, R., "More Instant Messaging Interoperability (MIMI) Identity Concepts", Work in Progress, Internet-Draft, draft-mahy-mimi-identity-00, 11 July 2022, <<https://www.ietf.org/archive/id/draft-mahy-mimi-identity-00.txt>>.
- [RFC2046] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, DOI 10.17487/RFC2046, November 1996, <<https://www.rfc-editor.org/info/rfc2046>>.
- [RFC3156] Elkins, M., Del Torto, D., Levien, R., and T. Roessler, "MIME Security with OpenPGP", RFC 3156, DOI 10.17487/RFC3156, August 2001, <<https://www.rfc-editor.org/info/rfc3156>>.

- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, DOI 10.17487/RFC5322, October 2008, <<https://www.rfc-editor.org/info/rfc5322>>.
- [RFC6120] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, DOI 10.17487/RFC6120, March 2011, <<https://www.rfc-editor.org/info/rfc6120>>.
- [RFC8098] Hansen, T., Ed. and A. Melnikov, Ed., "Message Disposition Notification", STD 85, RFC 8098, DOI 10.17487/RFC8098, February 2017, <<https://www.rfc-editor.org/info/rfc8098>>.
- [RFC8551] Schaad, J., Ramsdell, B., and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 4.0 Message Specification", RFC 8551, DOI 10.17487/RFC8551, April 2019, <<https://www.rfc-editor.org/info/rfc8551>>.
- [RFC9078] Crocker, D., Signes, R., and N. Freed, "Reaction: Indicating Summary Reaction to a Message", RFC 9078, DOI 10.17487/RFC9078, August 2021, <<https://www.rfc-editor.org/info/rfc9078>>.
- [W3C.CR-html52-20170808] Faulkner, S., Eicholz, A., Leithead, T., Danilo, A., and S. Moon, "HTML 5.2", World Wide Web Consortium CR CR-html52-20170808, 8 August 2017, <<https://www.w3.org/TR/2017/CR-html52-20170808>>.

Appendix A. TLS Presentation Language multipart container format

In a heterogenous group of IM clients, it is often desirable to send more than one media type as alternatives, such that IM clients have a choice of which media type to render. For example, imagine an IM group containing a set of clients which support a common video format and a subset which only support animated GIFs. The sender could send a multipart/alternative [RFC2046] container containing both media types. Every client in the group chat could render something resembling the media sent.

Likewise it is often desirable to send more than one media type intended to be rendered together as in (for example a rich text document with embedded images), which can be represented using the multipart/mixed [RFC2046] media type.

Some implementors complain that the multipart types are unnatural to use inside a binary protocol which requires explicit lengths such as MLS [I-D.ietf-mls-protocol]. Concretely, an implementation has to scan through the entire content to construct a boundary token which is not contained in the content.

While the author does not care about the specific syntax used, for comparison purposes presents a multipart container format using the TLS presentation language syntax used by the MLS protocol.

Note that there is a minor semantic difference between multipart/alternative and the proposal below. In multipart/alternative, the parts are presented in preference order by the sender. The receiver is support to render the first type which it supports. This container includes an ordering flag. As well, even if the flag is ordered, it is up to the IETF community to decide if it is acceptable for the receiver to choose its "best" format to render among an ordered preference list provided by the sender, or if the receiver must respect the ordered preference of the sender.

```
struct {
    /* a valid "Language-tag" as defined in RFC 5646 */
    opaque language_tag<1..52>;
} LanguageTag;

struct {
    ContentType content_type;
    LanguageTag content_languages<V>;
    opaque<V> body;
} Part;

enum {
    reserved(0),
    multipart_container_v1(1),
    (255)
} MultipartVersion;

enum {
    reserved(0),
    mixed(1),
    alternative(2),
    (255)
} MultipartSemantics;

enum {
    reserved(0),
    unordered(1),
    ordered(2),
    (255)
} MultipartOrdering;

struct {
    uint8 container_version;
    uint16 number_of_parts;
    MultipartSemantics semantics;
    MultipartOrdering ordering;
    Part parts<V>;
} MultipartContainer;
```

Author's Address

Rohan Mahy
Wire

Email: rohan.mahy@wire.com