

Network Working Group
Internet Draft

Nathaniel Borenstein
Ned Freed
<[draft-ietf-822ext-mime-imb-03.txt](#)>

Multipurpose Internet Mail Extensions
(MIME) Part One:

Format of Internet Message Bodies

May 5, 1995

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. Internet-Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet-Drafts as reference material or to cite them other than as a "working draft" or "work in progress".

To learn the current status of any Internet-Draft, please check the `1id-abstracts.txt` listing contained in the Internet-Drafts Shadow Directories on `ds.internic.net` (US East Coast), `nic.nordu.net` (Europe), `ftp.isi.edu` (US West Coast), or `munari.oz.au` (Pacific Rim).

1. Abstract

STD 11, [RFC 822](#), defines a message representation protocol specifying considerable detail about US-ASCII message headers, and leaves the message content, or message body, as flat US-ASCII text. This set of documents, collectively called the Multipurpose Internet Mail Extensions, or MIME, redefines the format of messages to allow for

- (1) textual message bodies in character sets other than US-ASCII,
- (2) non-textual message bodies,
- (3) multi-part message bodies, and
- (4) textual header information in character sets other than US-ASCII.

These documents are based on earlier work documented in RFC 934, STD 11, and [RFC 1049](#), but extends and revises them. Because [RFC 822](#) said so little about message bodies, these documents are largely orthogonal to (rather than a revision of) [RFC 822](#).

In particular, these documents are designed to provide facilities to include multiple parts in a single message, to represent body and header text in character sets other than US-ASCII, to represent formatted multi-font text messages, to represent non-textual material such as images and audio fragments, and generally to facilitate later extensions defining new types of Internet mail for use by cooperating mail agents.

This initial document specifies the various headers used to describe the structure of MIME messages. The second document, RFC MIME-INT, defines the general structure of the MIME media typing system and defines an initial set of media types. The third document, RFC MIME-HEADERS, describes extensions to [RFC 822 to allow non-US-ASCII text data in Internet mail header fields](#). The fourth document, RFC MIME-REG, specifies various IANA registration procedures for MIME-related entities. The fifth and final document, RFC MIME-CONF, describes MIME conformance criteria as well as providing some illustrative examples of MIME message formats, acknowledgements, and the bibliography.

These documents are revisions of RFCs 1521, 1522, and 1590, which themselves were revisions of RFCs 1341 and 1342. An appendix in RFC MIME-CONF describes differences and changes from previous versions.

Expires November 1995

[Page 2]

2. Table of Contents

1 Abstract	1
2 Table of Contents	3
3 Introduction	4
4 Definitions, Conventions, and Generic BNF Grammar	6
4.1 CRLF	7
4.2 Character Set	7
4.3 Message	8
4.4 Body Part	8
4.5 Entity	8
4.6 Body	8
4.7 7bit Data	8
4.8 8bit Data	9
4.9 Binary Data	9
4.10 Lines	9
5 MIME Header Fields	9
6 MIME-Version Header Field	10
7 Content-Type Header Field	12
7.1 Syntax of the Content-Type Header Field	14
7.2 Content-Type Defaults	16
8 Content-Transfer-Encoding Header Field	17
8.1 Content-Transfer-Encoding Syntax	17
8.2 Content-Transfer-Encodings Semantics	17
8.3 New Content-Transfer-Encodings	19
8.4 Interpretation and Use	19
8.5 Translating Encodings	21
8.6 Canonical Encoding Model	22
8.7 Quoted-Printable Content-Transfer-Encoding	22
8.8 Base64 Content-Transfer-Encoding	26
9 Content-ID Header Field	29
10 Content-Description Header Field	29
11 Additional MIME Header Fields	30
12 Summary	30
13 Security Considerations	30
14 Authors' Addresses	31
A Collected Grammar	32

Expires November 1995

[Page 3]

3. Introduction

Since its publication in 1982, [RFC 822](#) [[RFC-822](#)] has defined the standard format of textual mail messages on the Internet. Its success has been such that the [RFC 822](#) format has been adopted, wholly or partially, well beyond the confines of the Internet and the Internet SMTP transport defined by [RFC 821](#) [[RFC-821](#)]. As the format has seen wider use, a number of limitations have proven increasingly restrictive for the user community.

[RFC 822](#) was intended to specify a format for text messages. As such, non-text messages, such as multimedia messages that might include audio or images, are simply not mentioned. Even in the case of text, however, [RFC 822](#) is inadequate for the needs of mail users whose languages require the use of character sets richer than US-ASCII. Since [RFC 822](#) does not specify mechanisms for mail containing audio, video, Asian language text, or even text in most European languages, additional specifications are needed.

One of the notable limitations of [RFC 821/822](#) based mail systems is the fact that they limit the contents of electronic mail messages to relatively short lines (e.g. 1000 characters or less [[RFC821](#)]) of 7-bit US-ASCII. This forces users to convert any non-textual data that they may wish to send into seven-bit bytes representable as printable US-ASCII characters before invoking a local mail UA (User Agent, a program with which human users send and receive mail). Examples of such encodings currently used in the Internet include pure hexadecimal, uuencode, the 3-in-4 base 64 scheme specified in [RFC 1421](#), the Andrew Toolkit Representation [ATK], and many others.

The limitations of [RFC 822](#) mail become even more apparent as gateways are designed to allow for the exchange of mail messages between [RFC 822](#) hosts and X.400 hosts. X.400 [X400] specifies mechanisms for the inclusion of non-textual body parts within electronic mail messages. The current standards for the mapping of X.400 messages to [RFC 822](#) messages specify either that X.400 non-textual body parts must be converted to (not encoded in) IA5Text format, or that they must be discarded, notifying the [RFC 822](#) user that discarding has occurred. This is clearly undesirable, as information that a user may wish to receive is lost. Even though a user agent

Expires November 1995

[Page 4]

may not have the capability of dealing with the non-textual body part, the user might have some mechanism external to the UA that can extract useful information from the body part. Moreover, it does not allow for the fact that the message may eventually be gatewayed back into an X.400 message handling system (i.e., the X.400 message is "tunneled" through Internet mail), where the non-textual information would definitely become useful again.

This document describes several mechanisms that combine to solve most of these problems without introducing any serious incompatibilities with the existing world of [RFC 822](#) mail. In particular, it describes:

- (1) A MIME-Version header field, which uses a version number to declare a message to be conformant with this specification and allows mail processing agents to distinguish between such messages and those generated by older or non-conformant software, which are presumed to lack such a field.
- (2) A Content-Type header field, generalized from [RFC 1049](#) [[RFC-1049](#)], which can be used to specify the media type and subtype of data in the body of a message and to fully specify the native representation (canonical form) of such data.
- (3) A Content-Transfer-Encoding header field, which can be used to specify an auxiliary encoding that was applied to the data in order to allow it to pass through mail transport mechanisms which may have data or character set limitations.
- (4) Two additional header fields that can be used to further describe the data in a body, the Content-ID and Content-Description header fields.

All of the header fields defined in this document are subject to the general syntactic rules for header fields specified in [RFC 822](#). In particular, all of these header fields can include [RFC 822](#) comments, which have no semantic content and should be ignored during MIME processing.

Finally, to specify and promote interoperability, RFC MIME-CONF provides a basic applicability statement for a subset of

Expires November 1995

[Page 5]

the above mechanisms that defines a minimal level of "conformance" with this document.

HISTORICAL NOTE: Several of the mechanisms described in this document may seem somewhat strange or even baroque at first reading. It is important to note that compatibility with existing standards AND robustness across existing practice were two of the highest priorities of the working group that developed this document. In particular, compatibility was always favored over elegance.

Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol. [RFC 822](#) and [RFC 1123](#) [[RFC-1123](#)] also provide essential background for MIME since no conforming implementation of MIME can violate them. In addition, several other informational RFC documents will be of interest to the MIME implementor, in particular [RFC 1344](#) [[RFC-1344](#)], [RFC 1345](#) [[RFC-1345](#)], and [RFC 1524](#) [[RFC-1524](#)].

4. Definitions, Conventions, and Generic BNF Grammar

Although the mechanisms specified in this document are all described in prose, most are also described formally in the augmented BNF notation of [RFC 822](#). Implementors will need to be familiar with this notation in order to understand this specification, and are referred to [RFC 822](#) for a complete explanation of the augmented BNF notation.

Some of the augmented BNF in this document makes reference to syntactic entities that are defined in [RFC 822](#) and not in this document. A complete formal grammar, then, is obtained by combining [Appendix A](#) of this document, the collected grammar, with the BNF of [RFC 822](#) plus the modifications to [RFC 822](#) defined in [RFC 1123](#) (which specifically changes the syntax for 'return', 'date' and 'mailbox').

In this document, all numeric and octet values are given in decimal notation. All media type values, subtype values, and parameter names as defined in this document are case-insensitive. However, parameter values are case-sensitive unless otherwise specified for the specific parameter.

Expires November 1995

[Page 6]

FORMATTING NOTE: Notes, such as this one, provide additional nonessential information which may be skipped by the reader 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. Such information may in particular be skipped by those who are focused entirely on building a conformant implementation, but may be of use to those who wish to understand why certain design choices were made.

4.1. CRLF

The term CRLF, in this document, refers to the sequence of octets corresponding to the two US-ASCII characters CR (decimal value 13) and LF (decimal value 10) which, taken together, in this order, denote a line break in [RFC 822](#) mail.

4.2. Character Set

The term "character set" is used in this document to refer to a table-based method of converting a sequence of octets into a sequence of characters. Note that unconditional and unambiguous conversion in the other direction is not required, in that not all characters may be available in a given character set and a character set may provide more than one sequence of octets to represent a particular character. This definition is intended to allow various kinds of character encodings, from simple single-table mappings such as US-ASCII to complex table switching methods such as those that use ISO 2022's techniques. However, the definition associated with a MIME character set name must fully specify the mapping to be performed from octets to characters. In particular, use of external profiling information to determine the exact mapping is not permitted.

HISTORICAL NOTE: The term "character set" originated in the definition of US-ASCII and similar 7-bit and 8-bit specifications. These define true sets. However, the advent of multi-octet character encodings and switching techniques have transformed character sets into entities that properly speaking are no longer strictly sets. Some other communities have adopted the term "character encoding" for what MIME calls

Expires November 1995

[Page 7]

a "character set" as a result.

[4.3.](#) Message

The term "message", when not further qualified, means either the (complete or "top-level") message being transferred on a network, or a message encapsulated in a body part of type "message".

[4.4.](#) Body Part

The term "body part", in this document, refers to content headers and contents of either a message or one of the parts in the body of a multipart entity. A body part has a header and a body, so it makes sense to speak about the body of a body part.

[4.5.](#) Entity

The term "entity", in this document, means either a message or a body part. All kinds of entities share the property that they have a header and a body.

[4.6.](#) Body

The term "body", when not further qualified, means the body of an entity, that is the body of either a message or of a body part.

NOTE: The previous four definitions are clearly circular. This is unavoidable, since the overall structure of a MIME message is indeed recursive.

[4.7.](#) 7bit Data

"7bit data" refers to data that is all represented as relatively short lines with 998 octets or less between CRLF line separation sequences [[RFC821](#)]. No octets with decimal values greater than 127 are allowed and neither are NULs (octets with decimal value 0). CR (decimal value 13) and LF

Expires November 1995

[Page 8]

(decimal value 10) octets only occur as part of CRLF line separation sequences.

4.8. 8bit Data

"8bit data" refers to data that is all represented as relatively short lines with 998 octets or less between CRLF line separation sequences [[RFC821](#)]), but octets with decimal values greater than 127 may be used. As with "7bit data" CR and LF octets only occur as part of CRLF line separation sequences and no NULs are allowed.

4.9. Binary Data

"Binary data" refers to data where any sequence of octets whatsoever is allowed.

4.10. Lines

"Lines" are defined as sequences of octets separated by a CRLF sequences. This is consistent with both [RFC 821](#) and [RFC 822](#). "Lines" only refers to a unit of text in a message, which may or may not correspond to something that is actually displayed by a user agent.

5. MIME Header Fields

MIME defines a number of new [RFC 822](#) header fields that are used to describe the content of messages. These header fields occur in two contexts:

- (1) As part of a regular [RFC 822](#) message header.
- (2) In a MIME body part header within a multipart construct.

The formal definition of these header fields is as follows:

Expires November 1995

[Page 9]

```
MIME-message-headers := fields
                        version CRLF
                        [ content CRLF ]
                        [ encoding CRLF ]
                        [ id CRLF ]
                        [ description CRLF ]
                        *( mime-extension-field CRLF )
                        ; The ordering of the header
                        ; fields implied by this BNF
                        ; definition should be ignored
```

```
MIME-part-headers := [ content CRLF ]
                     [ encoding CRLF ]
                     [ id CRLF ]
                     [ description CRLF ]
                     *( mime-extension-field CRLF )
                     ; The ordering of the header
                     ; fields implied by this BNF
                     ; definition should be ignored
```

The syntax of the various specific MIME header fields will be described in the following sections.

6. MIME-Version Header Field

Since [RFC 822](#) was published in 1982, there has really been only one format standard for Internet messages, and there has been little perceived need to declare the format standard in use. This document is an independent document that complements [RFC 822](#). Although the extensions in this document have been defined in such a way as to be compatible with RFC 822, there are still circumstances in which it might be desirable for a mail-processing agent to know whether a message was composed with the new standard in mind.

Therefore, this document defines a new header field, "MIME-Version", which is to be used to declare the version of the Internet message body format standard in use.

Messages composed in accordance with this document **MUST** include such a header field, with the following verbatim text:

```
MIME-Version: 1.0
```

Expires November 1995

[Page 10]

The presence of this header field is an assertion that the message has been composed in compliance with this document.

Since it is possible that a future document might extend the message format standard again, a formal BNF is given for the content of the MIME-Version field:

```
version := "MIME-Version" ":" 1*DIGIT "." 1*DIGIT
```

Thus, future format specifiers, which might replace or extend "1.0", are constrained to be two integer fields, separated by a period. If a message is received with a MIME-version value other than "1.0", it cannot be assumed to conform with this specification.

Note that the MIME-Version header field is required at the top level of a message. It is not required for each body part of a multipart entity. It is required for the embedded headers of a body of type "message" if and only if the embedded message is itself claimed to be MIME-conformant.

It is not possible to fully specify how a mail reader that conforms with MIME as defined in this document should treat a message that might arrive in the future with some value of MIME-Version other than "1.0".

It is also worth noting that version control for specific media types is not accomplished using the MIME-Version mechanism. In particular, some formats (such as application/postscript) have version numbering conventions that are internal to the document format. Where such conventions exist, MIME does nothing to supersede them. Where no such conventions exist, a MIME media type might use a "version" parameter in the content-type field if necessary.

NOTE TO IMPLEMENTORS: When checking MIME-Version values any [RFC 822](#) comment strings that are present must be ignored. In particular, the following four MIME-Version fields are equivalent:

```
MIME-Version: 1.0
```

```
MIME-Version: 1.0 (produced by MetaSend Vx.x)
```

```
MIME-Version: (produced by MetaSend Vx.x) 1.0
```

Expires November 1995

[Page 11]

MIME-Version: 1.(produced by MetaSend Vx.x)0

In the absence of a MIME-Version field, a receiving user agent (whether MIME compliant or not) may optionally choose to interpret the body of the message according to local conventions. Many such conventions are currently in use and it should be noted that in practice non-MIME messages can contain just about anything.

It is impossible to be certain that a non-MIME message is actually plain text in the US-ASCII character set since it might well be a message that, using some set of nonstandard local conventions that predate this document, includes text in another character set or non-textual data presented in a manner that cannot be automatically recognized (e.g., a uuencoded compressed UNIX tar file).

MIME-compliant user agents are required, if they support any such nonstandard conventions at all, to do so on received messages only -- they must not send non-MIME messages containing anything other than US-ASCII text.

In particular, the use of non-US-ASCII text in messages without a MIME-Version field is strongly discouraged as it impedes interoperability when sending messages between regions with different localization conventions. MIME-compliant user agents **MUST** include proper MIME labelling when sending anything other than plain text in the US-ASCII character set.

In addition, non-MIME user agents should be upgraded if at all possible to include appropriate MIME header information in the messages they send even if nothing else in MIME is supported. This upgrade will have little, if any, effect on non-MIME recipients and will aid MIME in correctly displaying such messages. It also provides a smooth transition path to eventual adoption of other MIME capabilities.

7. Content-Type Header Field

The purpose of the Content-Type field is to describe the data contained in the body fully enough that the receiving user agent can pick an appropriate agent or mechanism to present the data to the user, or otherwise deal with the data in an appropriate manner. The value in this field is called a media

Expires November 1995

[Page 12]

type.

HISTORICAL NOTE: The Content-Type header field was first defined in [RFC 1049](#). [RFC 1049](#) used a simpler and less powerful syntax, but one that is largely compatible with the mechanism given here.

The Content-Type header field specifies the nature of the data in the body of an entity by giving media type and subtype identifiers, and by providing auxiliary information that may be required for certain media types. After the media type and subtype names, the remainder of the header field is simply a set of parameters, specified in an attribute=value notation. The ordering of parameters is not significant.

In general, the top-level media type is used to declare the general type of data, while the subtype specifies a specific format for that type of data. Thus, a media type of "image/xyz" is enough to tell a user agent that the data is an image, even if the user agent has no knowledge of the specific image format "xyz". Such information can be used, for example, to decide whether or not to show a user the raw data from an unrecognized subtype -- such an action might be reasonable for unrecognized subtypes of text, but not for unrecognized subtypes of image or audio. For this reason, registered subtypes of text, image, audio, and video should not contain embedded information that is really of a different type. Such compound formats should be represented using the "multipart" or "application" types.

Parameters are modifiers of the media subtype, and as such do not fundamentally affect the nature of the content. The set of meaningful parameters depends on the media type and subtype. Most parameters are associated with a single specific subtype. However, a given top-level media type may define parameters which are applicable to any subtype of that type. Parameters may be required by their defining content type or subtype or they may be optional. MIME implementations must ignore any parameters whose names they do not recognize.

For example, the "charset" parameter is applicable to any subtype of "text", while the "boundary" parameter is required for any subtype of the "multipart" media type.

Expires November 1995

[Page 13]

There are NO globally-meaningful parameters that apply to all media types. Truly global mechanisms are best addressed, in the MIME model, by the definition of additional Content-* header fields.

An initial set of seven top-level media types is defined in MIME-IMT. Five of these are discrete types whose content is essentially opaque as far as MIME processing is concerned. The remaining two are composite types whose contents require additional handling by MIME processors.

This set of top-level media types is intended to be substantially complete. It is expected that additions to the larger set of supported types can generally be accomplished by the creation of new subtypes of these initial types. In the future, more top-level types may be defined only by a standards-track extension to this standard. If another top-level type is to be used for any reason, it must be given a name starting with "X-" to indicate its non-standard status and to avoid a potential conflict with a future official name.

7.1. Syntax of the Content-Type Header Field

In the Augmented BNF notation of [RFC 822](#), a Content-Type header field value is defined as follows:

```
content := "Content-Type" ":" type "/" subtype
          *(";" parameter)
          ; Matching of media type and subtype
          ; is ALWAYS case-insensitive

type := discrete-type / composite-type

discrete-type := "text" / "image" / "audio" / "video" /
                 "application" / extension-token

composite-type := "message" / "multipart" / extension-token

extension-token := iana-token / ietf-token / x-token

iana-token := <a publicly-defined extension token,
               registered with IANA, as specified in
               RFC MIME-REG [RFC-MIME-REG]>
```

Expires November 1995

[Page 14]

ietf-token := <a publicly-defined extension token,
initially registered with IANA and
subsequently standardized by the IETF>

x-token := <The two characters "X-" or "x-" followed, with
no intervening white space, by any token>

subtype := extension-token

parameter := attribute "=" value

attribute := token

value := token / quoted-string

token := 1*<any (US-ASCII) CHAR except SPACE, CTLs,
or tspecials>

tspecials := "(" / ")" / "<" / ">" / "@" /
"," / ";" / ":" / "\" / <>
"/" / "[" / "]" / "?" / "="
; Must be in quoted-string,
; to use within parameter values

Note that the definition of "tspecials" is the same as the RFC [822](#) definition of "specials" with the addition of the three characters "/", "?", and "=", and the removal of ".".

Note also that a subtype specification is MANDATORY -- it may not be omitted from a Content-Type header field. As such, there are no default subtypes.

The type, subtype, and parameter names are not case sensitive. For example, TEXT, Text, and TeXt are all equivalent top-level media types. Parameter values are normally case sensitive, but sometimes are interpreted in a case-insensitive fashion, depending on the intended use. (For example, multipart boundaries are case-sensitive, but the "access-type" parameter for message/External-body is not case-sensitive.)

Note that the value of a quoted string parameter does not include the quotes. That is, the quotation marks in a quoted-string are not a part of the value of the parameter, but are merely used to delimit that parameter value. In addition, comments are allowed in accordance with [RFC 822](#)

Expires November 1995

[Page 15]

rules for structured header fields. Thus the following two forms

```
Content-type: text/plain; charset=us-ascii (Plain text)
```

```
Content-type: text/plain; charset="us-ascii"
```

are completely equivalent.

Beyond this syntax, the only syntactic constraint on the definition of subtype names is the desire that their uses must not conflict. That is, it would be undesirable to have two different communities using "Content-Type: application/foobar" to mean two different things. The process of defining new media subtypes, then, is not intended to be a mechanism for imposing restrictions, but simply a mechanism for publicizing their definition and usage. There are, therefore, two acceptable mechanisms for defining new media subtypes:

- (1) Private values (starting with "X-") may be defined bilaterally between two cooperating agents without outside registration or standardization.
- (2) New standard values **MUST** be documented, registered with, and approved by IANA, as described in RFC MIME-REG [RFC-MIME-REG].

The second document in this set, RFC MIME-IMT, defines the initial set of media types for MIME.

7.2. Content-Type Defaults

Default [RFC 822](#) messages without a MIME Content-Type header are taken by this protocol to be plain text in the US-ASCII character set, which can be explicitly specified as:

```
Content-type: text/plain; charset=us-ascii
```

This default is assumed if no Content-Type header field is specified. In the presence of a MIME-Version header field, a receiving User Agent can also assume that plain US-ASCII text was the sender's intent. Plain US-ASCII text may still be assumed in the absence of a MIME-Version specification, but the sender's intent might have been otherwise.

Expires November 1995

[Page 16]

8. Content-Transfer-Encoding Header Field

Many media types which could be usefully transported via email are represented, in their "natural" format, as 8-bit character or binary data. Such data cannot be transmitted over some transfer protocols. For example, [RFC 821](#) (SMTP) restricts mail messages to 7-bit US-ASCII data with lines no longer than **1000 characters including any trailing CRLF line separator**.

It is necessary, therefore, to define a standard mechanism for encoding such data into a 7-bit short line format. Proper labelling of unencoded material in less restrictive formats for direct use over less restrictive transports is also desirable. This document specifies that such encodings will be indicated by a new "Content-Transfer-Encoding" header field. This field has not been defined by any previous standard.

8.1. Content-Transfer-Encoding Syntax

The Content-Transfer-Encoding field's value is a single token specifying the type of encoding, as enumerated below. Formally:

```
encoding := "Content-Transfer-Encoding" ":" mechanism
```

```
mechanism := "7bit" / "8bit" / "binary" /  
             "quoted-printable" / "base64" /  
             ietf-token / x-token
```

These values are not case sensitive -- Base64 and BASE64 and bAsE64 are all equivalent. An encoding type of 7BIT requires that the body is already in a 7-bit mail-ready representation. This is the default value -- that is, "Content-Transfer-Encoding: 7BIT" is assumed if the Content-Transfer-Encoding header field is not present.

8.2. Content-Transfer-Encodings Semantics

This single Content-Transfer-Encoding token actually provides two pieces of information. It specifies what sort of encoding transformation the body was subjected to, and it specifies what the domain of the result is.

Expires November 1995

[Page 17]

Three transformations are currently defined: identity, the "quoted-printable" encoding, and the "base64" encoding. The domains are "binary", "8bit" and "7bit".

The Content-Transfer-Encoding values "7bit", "8bit", and "binary" all mean that the identity (i.e. NO) encoding transformation has been performed. As such, they serve simply as indicators of the domain of the body data, and provide useful information about the sort of encoding that might be needed for transmission in a given transport system. The terms "7bit data", "8bit data", and "binary data" are all defined in [Section 4](#).

The quoted-printable and base64 encodings transform their input from an arbitrary domain into material in the "7bit" range, thus making it safe to carry over restricted transports. The specific definition of the transformations are given below.

The proper Content-Transfer-Encoding label must always be used. Labelling unencoded data containing 8-bit characters as "7bit" is not allowed, nor is labelling unencoded non-line-oriented data as anything other than "binary" allowed.

Unlike media subtypes, a proliferation of Content-Transfer-Encoding values is both undesirable and unnecessary. However, establishing only a single transformation into the "7bit" domain does not seem possible. There is a tradeoff between the desire for a compact and efficient encoding of largely-binary data and the desire for a readable encoding of data that is mostly, but not entirely, 7-bit. For this reason, at least two encoding mechanisms are necessary: a "readable" encoding (quoted-printable) and a "dense" encoding (base64).

Mail transport for unencoded 8-bit data is defined in [RFC 1652](#) [[RFC-1652](#)]. As of the publication of this document, there are no standardized Internet mail transports for which it is legitimate to include unencoded binary data in mail bodies. Thus there are no circumstances in which the "binary" Content-Transfer-Encoding is actually valid in Internet mail. However, in the event that binary mail transport becomes a reality in Internet mail, or when this document is used in conjunction with any other binary-capable transport mechanism, binary bodies should be labelled as such using this mechanism.

Expires November 1995

[Page 18]

NOTE: The five values defined for the Content-Transfer-Encoding field imply nothing about the media type other than the algorithm by which it was encoded or the transport system requirements if unencoded.

8.3. New Content-Transfer-Encodings

Implementors may, if necessary, define private Content-Transfer-Encoding values, but must use an x-token, which is a name prefixed by "X-", to indicate its non-standard status, e.g., "Content-Transfer-Encoding: x-my-new-encoding". Additional standardized Content-Transfer-Encoding values must be specified by a standards-track RFC. Additional requirements such specifications must meet are given in RFC REG. As such, all content-transfer-encoding namespace except that beginning with "X-" is explicitly reserved to the IANA for future use.

Unlike media types and subtypes, the creation of new Content-Transfer-Encoding values is STRONGLY discouraged, as it seems likely to hinder interoperability with little potential benefit

8.4. Interpretation and Use

If a Content-Transfer-Encoding header field appears as part of a message header, it applies to the entire body of that message. If a Content-Transfer-Encoding header field appears as part of a body part's headers, it applies only to the body of that body part. If an entity is of type "multipart" the Content-Transfer-Encoding is not permitted to have any value other than "7bit", "8bit" or "binary". Even more severe restrictions apply to some subtypes of the "message" type.

It should be noted that most media types are defined in terms of octets rather than bits, so that the mechanisms described here are mechanisms for encoding arbitrary octet streams, not bit streams. If a bit stream is to be encoded via one of these mechanisms, it must first be converted to an 8-bit byte stream using the network standard bit order ("big-endian"), in which the earlier bits in a stream become the higher-order bits in a 8-bit byte. A bit stream not ending at an 8-bit

Expires November 1995

[Page 19]

boundary must be padded with zeroes. This document provides a mechanism for noting the addition of such padding in the case of the application/octet-stream media type, which has a "padding" parameter.

The encoding mechanisms defined here explicitly encode all data in US-ASCII. Thus, for example, suppose an entity has header fields such as:

```
Content-Type: text/plain; charset=ISO-8859-1
Content-transfer-encoding: base64
```

This must be interpreted to mean that the body is a base64 US-ASCII encoding of data that was originally in ISO-8859-1, and will be in that character set again after decoding.

Certain Content-Transfer-Encoding values may only be used on certain media types. In particular, it is EXPRESSLY FORBIDDEN to use any encodings other than "7bit", "8bit", or "binary" with any composite media type, i.e. one that recursively includes other Content-Type fields. Currently the only composite media types are "multipart" and "message". All encodings that are desired for bodies of type multipart or message must be done at the innermost level, by encoding the actual body that needs to be encoded.

It should also be noted that, by definition, if a composite entity has a transfer-encoding value such as "7bit", but one of the enclosed parts has a less restrictive value such as "8bit", then either the outer "7bit" labelling is in error, because 8-bit data are included, or the inner "8bit" labelling placed an unnecessarily high demand on the transport system because the actual included data were actually 7-bit-safe.

NOTE ON ENCODING RESTRICTIONS: Though the prohibition against using content-transfer-encodings on composite body data may seem overly restrictive, it is necessary to prevent nested encodings, in which data are passed through an encoding algorithm multiple times, and must be decoded multiple times in order to be properly viewed. Nested encodings add considerable complexity to user agents: Aside from the obvious efficiency problems with such multiple encodings, they can obscure the basic structure of a message. In particular, they can imply that several decoding operations are necessary simply to find out what types of bodies a message contains.

Expires November 1995

[Page 20]

Banning nested encodings may complicate the job of certain mail gateways, but this seems less of a problem than the effect of nested encodings on user agents.

NOTE ON THE RELATIONSHIP BETWEEN CONTENT-TYPE AND CONTENT-TRANSFER-ENCODING: It may seem that the Content-Transfer-Encoding could be inferred from the characteristics of the media that is to be encoded, or, at the very least, that certain Content-Transfer-Encodings could be mandated for use with specific media types. There are several reasons why this is not the case. First, given the varying types of transports used for mail, some encodings may be appropriate for some combinations of media types and transports but not for others. (For example, in an 8-bit transport, no encoding would be required for text in certain character sets, while such encodings are clearly required for 7-bit SMTP.)

Second, certain media types may require different types of transfer encoding under different circumstances. For example, many PostScript bodies might consist entirely of short lines of 7-bit data and hence require no encoding at all. Other PostScript bodies (especially those using Level 2 PostScript's binary encoding mechanism) may only be reasonably represented using a binary transport encoding. Finally, since the Content-Type field is intended to be an open-ended specification mechanism, strict specification of an association between media types and encodings effectively couples the specification of an application protocol with a specific lower-level transport. This is not desirable since the developers of a media type should not have to be aware of all the transports in use and what their limitations are.

8.5. Translating Encodings

The quoted-printable and base64 encodings are designed so that conversion between them is possible. The only issue that arises in such a conversion is the handling of line breaks. When converting from quoted-printable to base64 a line break must be converted into a CRLF sequence. Similarly, a CRLF sequence in base64 data must be converted to a quoted-printable line break, but ONLY when converting text data.

Expires November 1995

[Page 21]

8.6. Canonical Encoding Model

There was some confusion, in the predecessors of this RFC, regarding the model for when email data was to be converted to canonical form and encoded, and in particular how this process would affect the treatment of CRLFs, given that the representation of newlines varies greatly from system to system, and the relationship between content-transfer-encodings and character sets. A canonical model for encoding is presented in RFC MIME-CONF for this reason.

8.7. Quoted-Printable Content-Transfer-Encoding

The Quoted-Printable encoding is intended to represent data that largely consists of octets that correspond to printable characters in the US-ASCII character set. It encodes the data in such a way that the resulting octets are unlikely to be modified by mail transport. If the data being encoded are mostly US-ASCII text, the encoded form of the data remains largely recognizable by humans. A body which is entirely US-ASCII may also be encoded in Quoted-Printable to ensure the integrity of the data should the message pass through a character-translating, and/or line-wrapping gateway.

In this encoding, octets are to be represented as determined by the following rules:

- (1) (General 8-bit representation) Any octet, except a CR or LF that is part of a CRLF line break of the canonical (standard) form of the data being encoded, may be represented by an "=" followed by a two digit hexadecimal representation of the octet's value. The digits of the hexadecimal alphabet, for this purpose, are "0123456789ABCDEF". Uppercase letters must be used when sending hexadecimal data, though a robust implementation may choose to recognize lowercase letters on receipt. Thus, for example, the decimal value 12 (US-ASCII form feed) can be represented by "=0C", and the decimal value 61 (US-ASCII EQUAL SIGN) can be represented by "=3D". This rule must be followed except when the following rules allow an alternative encoding.

Expires November 1995

[Page 22]

- (2) (Literal representation) Octets with decimal values of 33 through 60 inclusive, and 62 through 126, inclusive, MAY be represented as the US-ASCII characters which correspond to those octets (EXCLAMATION POINT through LESS THAN, and GREATER THAN through TILDE, respectively).
- (3) (White Space) Octets with values of 9 and 32 MAY be represented as US-ASCII TAB (HT) and SPACE characters, respectively, but MUST NOT be so represented at the end of an encoded line. Any TAB (HT) or SPACE characters on an encoded line MUST thus be followed on that line by a printable character. In particular, an "=" at the end of an encoded line, indicating a soft line break (see rule #5) may follow one or more TAB (HT) or SPACE characters. It follows that an octet with decimal value 9 or 32 appearing at the end of an encoded line must be represented according to Rule #1. This rule is necessary because some MTAs (Message Transport Agents, programs which transport messages from one user to another, or perform a part of such transfers) are known to pad lines of text with SPACES, and others are known to remove "white space" characters from the end of a line. Therefore, when decoding a Quoted-Printable body, any trailing white space on a line must be deleted, as it will necessarily have been added by intermediate transport agents.
- (4) (Line Breaks) A line break in a text body, represented as a CRLF sequence in the text canonical form, must be represented by a ([RFC 822](#)) line break, which is also a CRLF sequence, in the Quoted-Printable encoding. Since the canonical representation of media types other than text do not generally include the representation of line breaks as CRLF sequences, no hard line breaks (i.e. line breaks that are intended to be meaningful and to be displayed to the user) should occur in the quoted-printable encoding of such types. Sequences like "=0D", "=0A", "=0A=0D" and "=0D=0A" will routinely appear in non-text data represented in quoted-printable, of course.

Note that many implementations may elect to encode the local representation of various content types directly rather than converting to canonical form first,

Expires November 1995

[Page 23]

encoding, and then converting back to local representation. In particular, this may apply to plain text material on systems that use newline conventions other than a CRLF terminator sequence. Such an implementation optimization is permissible, but only when the combined canonicalization-encoding step is equivalent to performing the three steps separately.

- (5) (Soft Line Breaks) The Quoted-Printable encoding REQUIRES that encoded lines be no more than 76 characters long. If longer lines are to be encoded with the Quoted-Printable encoding, "soft" line breaks must be used. An equal sign as the last character on a encoded line indicates such a non-significant ("soft") line break in the encoded text.

Thus if the "raw" form of the line is a single unencoded line that says:

Now's the time for all folk to come to the aid of their country.

This can be represented, in the Quoted-Printable encoding, as:

```
Now's the time =  
for all folk to come=  
to the aid of their country.
```

This provides a mechanism with which long lines are encoded in such a way as to be restored by the user agent. The 76 character limit does not count the trailing CRLF, but counts all other characters, including any equal signs.

Since the hyphen character ("-") is represented as itself in the Quoted-Printable encoding, care must be taken, when encapsulating a quoted-printable encoded body inside one or more multipart entities, to ensure that the boundary delimiter does not appear anywhere in the encoded body. (A good strategy is to choose a boundary that includes a character sequence such as "=_ " which can never appear in a quoted-printable body. See the definition of multipart messages in MIME-IMT.)

NOTE: The quoted-printable encoding represents something of a compromise between readability and reliability in transport. Bodies encoded with the quoted-printable encoding will work

Expires November 1995

[Page 24]

reliably over most mail gateways, but may not work perfectly over a few gateways, notably those involving translation into EBCDIC. A higher level of confidence is offered by the base64 Content-Transfer-Encoding. A way to get reasonably reliable transport through EBCDIC gateways is to also quote the US-ASCII characters

!"#\$%&'\]^_`{|}~

according to rule #1.

Because quoted-printable data is generally assumed to be line-oriented, it is to be expected that the representation of the breaks between the lines of quoted printable data may be altered in transport, in the same manner that plain text mail has always been altered in Internet mail when passing between systems with differing newline conventions. If such alterations are likely to constitute a corruption of the data, it is probably more sensible to use the base64 encoding rather than the quoted-printable encoding.

WARNING TO IMPLEMENTORS: If binary data are encoded in quoted-printable, care must be taken to encode CR and LF characters as "=0D" and "=0A", respectively. In particular, a CRLF sequence in binary data should be encoded as "=0D=0A". Otherwise, if CRLF were represented as a hard line break, it might be incorrectly decoded on platforms with different line break conventions.

For formalists, the syntax of quoted-printable data is described by the following grammar:

```
quoted-printable := qp-line *(CRLF qp-line)

qp-line := *(qp-segment transport-padding CRLF)
          qp-part transport-padding

qp-part := qp-section
          ; Maximum length of 76 characters

qp-segment := qp-section *(SPACE / TAB) "="
             ; Maximum length of 76 characters

qp-section := [*(ptext / SPACE / TAB) ptext]
```


Expires November 1995

[Page 25]

ptext := octet / safe-char

safe-char := <any octet with decimal value of 33 through
60 inclusive, and 62 through 126>
; Characters not listed as "mail-safe" in
; RFC MIME-CONF are also not recommended.

octet := "=" 2(DIGIT / "A" / "B" / "C" / "D" / "E" / "F")
; Octet must be used for characters > 127, =,
; SPACES or TABs at the ends of lines, and is
; recommended for any character not listed in
; RFC MIME-CONF as "mail-safe".

transport-padding := *LWSP-char
; Composers MUST NOT generate
; non-zero length transport
; padding, but receivers MUST
; be able to handle padding
; added by message transports.

IMPORTANT NOTE: The addition of LWSP between the elements
shown in this BNF is NOT allowed since this BNF does not
specify a structured header field.

8.8. Base64 Content-Transfer-Encoding

The Base64 Content-Transfer-Encoding is designed to represent
arbitrary sequences of octets in a form that need not be
humanly readable. The encoding and decoding algorithms are
simple, but the encoded data are consistently only about 33
percent larger than the unencoded data. This encoding is
virtually identical to the one used in Privacy Enhanced Mail
(PEM) applications, as defined in [RFC 1421](#) [[RFC-1421](#)].

A 65-character subset of US-ASCII is used, enabling 6 bits to
be represented per printable character. (The extra 65th
character, "=", is used to signify a special processing
function.)

NOTE: This subset has the important property that it is
represented identically in all versions of ISO 646, including
US-ASCII, and all characters in the subset are also
represented identically in all versions of EBCDIC. Other
popular encodings, such as the encoding used by the uuencode

Expires November 1995

[Page 26]

utility and the base85 encoding specified as part of Level 2 PostScript, do not share these properties, and thus do not fulfill the portability requirements a binary transport encoding for mail must meet.

The encoding process represents 24-bit groups of input bits as output strings of 4 encoded characters. Proceeding from left to right, a 24-bit input group is formed by concatenating 3 8-bit input groups. These 24 bits are then treated as 4 concatenated 6-bit groups, each of which is translated into a single digit in the base64 alphabet. When encoding a bit stream via the base64 encoding, the bit stream must be presumed to be ordered with the most-significant-bit first. That is, the first bit in the stream will be the high-order bit in the first 8-bit byte, and the eighth bit will be the low-order bit in the first 8-bit byte, and so on.

Each 6-bit group is used as an index into an array of 64 printable characters. The character referenced by the index is placed in the output string. These characters, identified in Table 1, below, are selected so as to be universally representable, and the set excludes characters with particular significance to SMTP (e.g., ".", CR, LF) and to the multipart boundary delimiters defined in MIME-IMT (e.g., "-").

Table 1: The Base64 Alphabet

Value	Encoding	Value	Encoding	Value	Encoding	Value	Encoding
0	A	17	R	34	i	51	z
1	B	18	S	35	j	52	0
2	C	19	T	36	k	53	1
3	D	20	U	37	l	54	2
4	E	21	V	38	m	55	3
5	F	22	W	39	n	56	4
6	G	23	X	40	o	57	5
7	H	24	Y	41	p	58	6
8	I	25	Z	42	q	59	7
9	J	26	a	43	r	60	8
10	K	27	b	44	s	61	9
11	L	28	c	45	t	62	+
12	M	29	d	46	u	63	/
13	N	30	e	47	v		
14	O	31	f	48	w	(pad)	=
15	P	32	g	49	x		
16	Q	33	h	50	y		

Expires November 1995

[Page 27]

The encoded output stream must be represented in lines of no more than 76 characters each. All line breaks or other characters not found in Table 1 must be ignored by decoding software. In base64 data, characters other than those in Table 1, line breaks, and other white space probably indicate a transmission error, about which a warning message or even a message rejection might be appropriate under some circumstances.

Special processing is performed if fewer than 24 bits are available at the end of the data being encoded. A full encoding quantum is always completed at the end of a body. When fewer than 24 input bits are available in an input group, zero bits are added (on the right) to form an integral number of 6-bit groups. Padding at the end of the data is performed using the "=" character. Since all base64 input is an integral number of octets, only the following cases can arise: (1) the final quantum of encoding input is an integral multiple of 24 bits; here, the final unit of encoded output will be an integral multiple of 4 characters with no "=" padding, (2) the final quantum of encoding input is exactly 8 bits; here, the final unit of encoded output will be two characters followed by two "=" padding characters, or (3) the final quantum of encoding input is exactly 16 bits; here, the final unit of encoded output will be three characters followed by one "=" padding character.

Because it is used only for padding at the end of the data, the occurrence of any "=" characters may be taken as evidence that the end of the data has been reached (without truncation in transit). No such assurance is possible, however, when the number of octets transmitted was a multiple of three and no "=" characters are present.

Any characters outside of the base64 alphabet are to be ignored in base64-encoded data.

Care must be taken to use the proper octets for line breaks if base64 encoding is applied directly to text material that has not been converted to canonical form. In particular, text line breaks must be converted into CRLF sequences prior to base64 encoding. The important thing to note is that this may be done directly by the encoder rather than in a prior canonicalization step in some implementations.

Expires November 1995

[Page 28]

NOTE: There is no need to worry about quoting potential boundary delimiters within base64-encoded parts of multipart entities because no hyphen characters are used in the base64 encoding.

9. Content-ID Header Field

In constructing a high-level user agent, it may be desirable to allow one body to make reference to another. Accordingly, bodies may be labelled using the "Content-ID" header field, which is syntactically identical to the "Message-ID" header field:

```
id := "Content-ID" ":" msg-id
```

Like the Message-ID values, Content-ID values must be generated to be world-unique.

The Content-ID value may be used for uniquely identifying MIME entities in several contexts, particularly for caching data referenced by the message/external-body mechanism. Although the Content-ID header is generally optional, its use is MANDATORY in implementations which generate data of the optional MIME media type "message/external-body". That is, each message/external-body entity must have a Content-ID field to permit caching of such data.

It is also worth noting that the Content-ID value has special semantics in the case of the multipart/alternative media type. This is explained in the section of MIME-IMT dealing with multipart/alternative.

10. Content-Description Header Field

The ability to associate some descriptive information with a given body is often desirable. For example, it may be useful to mark an "image" body as "a picture of the Space Shuttle Endeavor." Such text may be placed in the Content-Description header field. This header field is always optional.

```
description := "Content-Description" ":" *text
```

The description is presumed to be given in the US-ASCII

Expires November 1995

[Page 29]

character set, although the mechanism specified in RFC MIME-HEADERS [RFC-MIME-HEADERS] may be used for non-US-ASCII Content-Description values.

11. Additional MIME Header Fields

Future documents may elect to define additional MIME header fields for various purposes. Any new header field that further describes the content of a message should begin with the string "Content-" to allow such fields which appear in a message header to be distinguished from ordinary [RFC 822](#) message header fields.

MIME-extension-field := <Any [RFC 822](#) header field which
begins with the string
"Content-">

12. Summary

Using the MIME-Version, Content-Type, and Content-Transfer-Encoding header fields, it is possible to include, in a standardized way, arbitrary types of data objects with [RFC 822](#) conformant mail messages. No restrictions imposed by either [RFC 821](#) or [RFC 822](#) are violated, and care has been taken to avoid problems caused by additional restrictions imposed by the characteristics of some Internet mail transport mechanisms (see RFC MIME-CONF).

The next document in this set, RFC MIME-IMT, specifies the media types that can be labelled and transported using these headers.

13. Security Considerations

Security issues are discussed in the second document in this set, RFC MIME-IMT.

Expires November 1995

[Page 30]

14. Authors' Addresses

For more information, the authors of this document are best contacted via Internet mail:

Nathaniel S. Borenstein
First Virtual Holdings
25 Washington Avenue
Morristown, NJ 07960
USA

Email: nsb@nsb.fv.com
Phone: +1 201 540 8967
Fax: +1 201 993 3032

Ned Freed
Innosoft International, Inc.
1050 East Garvey Avenue South
West Covina, CA 91790
USA

Email: ned@innosoft.com
Phone: +1 818 919 3600
Fax: +1 818 919 3614

MIME is a result of the work of the Internet Engineering Task Force Working Group on Email Extensions. The chairman of that group, Greg Vaudreuil, may be reached at:

Gregory M. Vaudreuil
Tigon Corporation
17060 Dallas Parkway
Dallas Texas, 75248

Email: greg.vaudreuil@ons.octel.com
Phone: +1 214 733 2722

Expires November 1995

[Page 31]

[Appendix A](#) -- Collected Grammar

This appendix contains the complete BNF grammar for all the syntax specified by this document.

By itself, however, this grammar is incomplete. It refers to several entities that are defined by [RFC 822](#). Rather than reproduce those definitions here, and risk unintentional differences between the two, this document simply refers the reader to [RFC 822](#) for the remaining definitions. Wherever a term is undefined, it refers to the [RFC 822](#) definition.

```
attribute := token
```

```
composite-type := "message" / "multipart" / extension-token
```

```
content := "Content-Type" ":" type "/" subtype
          *(";" parameter)
          ; Matching of media type and subtype
          ; is ALWAYS case-insensitive
```

```
description := "Content-Description" ":" *text
```

```
discrete-type := "text" / "image" / "audio" / "video" /
                 "application" / extension-token
```

```
encoding := "Content-Transfer-Encoding" ":" mechanism
```

```
extension-token := iana-token / ietf-token / x-token
```

```
iana-token := <a publicly-defined extension token,
               registered with IANA, as specified in
               RFC MIME-REG>
```

```
ietf-token := <a publicly-defined extension token,
               initially registered with IANA and
               subsequently standardized by the IETF>
```

```
id := "Content-ID" ":" msg-id
```

Expires November 1995

[Page 32]

```
mechanism := "7bit" / "8bit" / "binary" /
             "quoted-printable" / "base64" /
             ietf-token / x-token

octet := "=" 2(DIGIT / "A" / "B" / "C" / "D" / "E" / "F")
        ; Octet must be used for characters > 127, =,
        ; SPACES or TABs at the ends of lines, and is
        ; recommended for any character not listed in
        ; RFC MIME-CONF as "mail-safe".

parameter := attribute "=" value

ptext := octet / safe-char

qp-line := *(qp-segment transport-padding CRLF)
          qp-part transport-padding

qp-part := qp-section
          ; Maximum length of 76 characters

qp-section := [*(ptext / SPACE / TAB) ptext]

qp-segment := qp-section *(SPACE / TAB) "="
             ; Maximum length of 76 characters

quoted-printable := qp-line *(CRLF qp-line)

safe-char := <any octet with decimal value of 33 through
             60 inclusive, and 62 through 126>
             ; Characters not listed as "mail-safe" in
             ; RFC MIME-CONF are also not recommended.

subtype := extension-token

token := 1*<any (US-ASCII) CHAR except SPACE, CTLs,
        or tspecials>

transport-padding := *LWSP-char
                  ; Composers MUST NOT generate
                  ; non-zero length transport
                  ; padding, but receivers MUST
                  ; be able to handle padding
                  ; added by message transports.
```


Expires November 1995

[Page 33]

```
tspecials := "(" / ")" / "<" / ">" / "@" /  
              "," / ";" / ":" / "\" / "<" / ">" /  
              "/" / "[" / "]" / "?" / "="  
              ; Must be in quoted-string,  
              ; to use within parameter values
```

```
type := discrete-type / composite-type
```

```
value := token / quoted-string
```

```
version := "MIME-Version" ":" 1*DIGIT "." 1*DIGIT
```

```
x-token := <The two characters "X-" or "x-" followed, with  
           no intervening white space, by any token>
```