

Implications of MIME for Internet Mail Gateways

Status of This Memo

This is an informational memo for the Internet community, and requests discussion and suggestions for improvements. This memo does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

The recent development of MIME (Multipurpose Internet Mail Extensions) offers a wide range of new opportunities for electronic mail system systems. Most of these opportunities are relevant only to user agents, the programs that interact with human users when they send and receive mail. However, some opportunities are also opened up for mail transport systems. While MIME was carefully designed so that it does not require any changes to Internet electronic message transport facilities, there are several ways in which message transport systems may want to take advantage of MIME. These opportunities are the subject of this memo.

Background -- The MIME Format

Recently, a new standardized format has been defined for enhanced electronic mail messages on the Internet. This format, known as MIME, permits messages to include, in a standardized manner, non-ASCII text, images, audio, and a variety of other kinds of interesting data.

The MIME effort was explicitly focused on requiring absolutely no changes at the message transport level. Because of this fact, MIME-format mail runs transparently on all known Internet or Internet-style mail systems. This means that those concerned solely with the maintenance and development of message transport services can safely ignore MIME completely, if they so choose.

However, the fact that MIME can be ignored, for the purpose of message transport, does not necessarily mean that it should be ignored. In particular, MIME offers several features that should be of interest to those responsible for message transport services. By exploiting these features, transport systems can provide certain additional kinds of service that are currently unavailable, and can alleviate a few existing problems.

The remainder of this document is an attempt to briefly point out and summarize some important ways in which MIME may be of use for message transport systems. This document makes no attempt to present a complete technical description of MIME, however. For that, the reader is referred to the MIME document itself [RFC-1341].

Mail Transport and Gateway Services: A Key Distinction

Before implementing any of the mechanisms discussed in this memo, one should be familiar with the distinction between mail transport service and mail gateway service. Basically, mail transport software is responsible for moving a message within a homogeneous electronic mail service network. Mail gateways, on the other hand, exchange mail between two significantly different mail environments, including via non-electronic services, such as postal mail.

In general, it is widely considered unacceptable for mail transport services to alter the contents of messages. In the case of mail gateways, however, such alteration is often inevitable. Thus, strictly speaking, many of the mechanisms described here apply only to gateways, and should not be used in simple mail transport systems. However, it is possible that some very special situations -- e.g., an SMTP relay that transports mail across extremely expensive intercontinental network links -- might need to modify messages, in order to provide appropriate service for those situations, and hence must redefine its role to be that of a gateway.

In this memo, it is assumed that transformations which alter a message's contents will be performed only by gateways, but it is recognized that some existing mail transport agents may choose to reclassify themselves as gateways in order to perform the functions described here.

Rejected Messages

An unfortunately frequent duty of message transport services is the rejection of mail to the sender. This may happen because the mail was undeliverable, or because it did not conform to the requirements of a gateway (e.g., it was too large).

There has never been a standard format for rejected messages in the past. This has been an annoyance, but not a major problem for text messages. For non-text messages, however, the lack of a standard rejection format is more crucial, because rejected messages typically appear to be text, and the user who finds himself viewing images or audio as if they were text is rarely happy with the result.

MIME makes it very easy to encapsulate messages in such a way that their semantics are completely preserved. The simplest way to do this is to make each rejection notice a MIME "multipart/mixed" message. That multipart message would contain two parts, a text part explaining the reason for the rejection, and an encapsulated message part that contained the rejected message itself.

It should be stressed that the transport software does not need to understand the structure of the rejected message at all. It merely needs to encapsulate it properly. The following, for example, shows how any MIME message may be encapsulated in a rejection message in such a way that all information will be immediately visible in the correct form if the recipient reads it with a MIME-conformant mail reader:

```
From: Mailer-Daemon <daemon@somewhere.com>
Subject: Rejected Message
Content-type: multipart/mixed; boundary=unique-boundary
```

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

A mail message you sent was rejected. The details of the rejected message are as follows:

```
From: Nathaniel Borenstein <nsb@bellcore.com>
Message-ID: <12345@bellcore.com>
To: bush@whitehouse.gov
Subject: I know my rights!
Rejection-reason: No mail from libertarians is accepted.
```

The original message follows below.

```
--unique-boundary
Content-type: message/rfc822
```

The ENTIRE REJECTED MESSAGE, starting with the headers, goes here.

```
--unique-boundary--
```

In the above example, the **ONLY** thing that is not "boilerplate" is the choice of boundary string. The phrase "unique-boundary" should be replaced by a string that does not appear (prefixed by two hyphens) in any of the body parts.

Encapsulating a message in this manner is very easily done, and will constitute a significant service that message transport services can perform for MIME users.

IMPORTANT NOTE: The format given above is simply one of many possible ways to format a rejection message using MIME. Independent IETF efforts are needed in order to standardize the format of rejections and acknowledgements.

Fragmenting and Reassembling Large Messages

One problem that occurs with increasing frequency in Internet mail is the rejection of messages because of size limitations. This problem can be expected to grow substantially more severe with the acceptance of MIME, as MIME invites the use of very large objects such as images and audio clips. Fortunately, MIME also provides mechanisms that can help alleviate the problem.

One particularly relevant MIME type is "message/partial", which can be used for the automatic fragmentation and reassembly of large mail messages. The message/partial type can be handled entirely at the user agent level, but message transport services can also make use of this type to provide more intelligent behavior at gateways.

In particular, when gatewaying mail to or from a system or network known to enforce size limitations that are more or less stringent than are enforced locally, message transport services might choose either to break a large message into fragments, or (perhaps less likely) to reassemble fragments into a larger message. The combination of these two behaviors can make the overall Internet mail environment appear more complete and seamless than it actually is.

Details on the message/partial format may be found in the MIME document. What follows is an example of how a simple short message might be broken into two message/partial messages. In practice, of course, the message/partial facility would only be likely to be used for much longer messages.

The following initial message:

```
From: Nathaniel Borenstein <nsb@bellcore.com>
To: Ned Freed: <ned@innosoft.com>
Subject: a test message
Content-type: image/gif
Content-Transfer-Encoding: base64
```

```
R0lGODdhQAGMabMAAAAAAP/u7swzIu6ZiLsiEdlEM+5VRGaI3WYAAO67qkRV
uwARd6q7/ywAAAAAQAGMAUME/hDISau900vNu/9gKI6kRJwoUa5s675wLM90l
XW5YKxqPyKRYgxv2dr4czwlMCZrQLFTYHBJ2hlyQYFiaz+i0WWBou7fOq1x8vXWfU
qU1fJ2qEhYaHGjhZQmJ2QT1xBW1ak1xUdV0/VjtsbpUEDaEJCQOIpqeoNV+LXo5W
fVN3dZKceAQPvgyhwQ2lqcXGxx5wja59eJIGUNCszF90sYp50CoqFZ4DoqMMo6M
```

can be transformed, invertibly, into the following two message/partial messages:

```
From: Nathaniel Borenstein <nsb@bellcore.com>
To: Ned Freed <ned@innosoft.com>
Subject: a test message
Content-type: message/partial; id="xyx@host.com";
        number=1; total=2
```

```
Content-type: image/gif
Content-Transfer-Encoding: base64
```

```
R0lGODdhQAGMabMAAAAAAP/u7swzIu6ZiLsiEdlEM+5VRGaI3WYAAO67qkRV
```

and

```
From: Nathaniel Borenstein <nsb@bellcore.com>
To: Ned Freed <ned@innosoft.com>
Subject: a test message
Content-type: message/partial; id="xyx@host.com";
        number=2; total=2
```

```
uwARd6q7/ywAAAAAQAGMAUME/hDISau900vNu/9gKI6kRJwoUa5s675wLM90l
XW5YKxqPyKRYgxv2dr4czwlMCZrQLFTYHBJ2hlyQYFiaz+i0WWBou7fOq1x8vXWfU
qU1fJ2qEhYaHGjhZQmJ2QT1xBW1ak1xUdV0/VjtsbpUEDaEJCQOIpqeoNV+LXo5W
fVN3dZKceAQPvgyhwQ2lqcXGxx5wja59eJIGUNCszF90sYp50CoqFZ4DoqMMo6M
```

Fragmenting such messages rather than rejecting them might be a reasonable option for some gateway services, at least for a certain range of message sizes. Of course, it is often difficult for a gateway to know what size limitations will be encountered "downstream", but intelligent guesses are often possible. Moreover, an IETF working group on SMTP extensions has proposed augmenting SMTP with a "SIZE" verb that would facilitate this process, thereby possibly requiring fragmentation on the fly during message

transmission.

Note also that fragmentation or reassembly might reasonably be performed, in differing circumstances, by either the sending or receiving gateway systems, depending on which system knew more about the capabilities of the other.

Using or Removing External-Body Pointers

Another MIME type oriented to extremely large messages is the "message/external-body" type. In this type of message, all or part of the body data is not included in the actual message itself. Instead, the Content-Type header field includes information that tells how the body data can be retrieved -- either via a file system, via anonymous ftp, or via other mechanisms.

The message/external-body type provides a new option for mail transport services that wishes to optimize the way bandwidth resources are used in a given environment. For example, the basic use of message/external-body is to reduce bandwidth in email traffic. However, when email crosses a slow and expensive boundary -- e.g., a satellite link across the Pacific -- it might make sense to retrieve the data itself and transform the external-body reference into the actual data. Alternately, it might make sense to copy the data itself to a new location, closer to the message recipients, and change the location pointed to in the message. Because the external-body specification can include an expiration date, message transport services can trade off storage and bandwidth capabilities to try to optimize the overall use of resources for very large messages.

Such behaviors by a gateway require careful analysis of cost/benefit tradeoffs and would be a dramatic departure from typical mail transport services. However, the potential benefits are quite significant, so that such the appropriate use of these service options should be explored.

For example, the following message includes PostScript data by external reference:

```
From: Nathaniel Borenstein <nsb@bellcore.com>
To: Ned Freed <ned@innosoft.com>
Subject: The latest MIME draft
Content-Type: message/external-body;
    name="BodyFormats.ps";
    site="thumper.bellcore.com";
    access-type=ANON-FTP;
    directory="pub";
    mode="image";
    expiration="Fri, 14 Jun 1991 19:13:14 -0400 (EDT)"

Content-type: application/postscript
```

A gateway to Australia might choose to copy the file to an Australian FTP archive, changing the relevant parameters on the Content-type header field. Alternately, it might choose simply to transform the message into one in which all the data were included:

```
From: Nathaniel Borenstein <nsb@bellcore.com>
To: Ned Freed <ned@innosoft.com>
Subject: The latest MIME draft
Content-type: application/postscript
```

```
%!PS-Adobe-1.0
%%Creator: greenbush:nsb (Nathaniel Borenstein,MRE 2A-
274,4270,9938586,21462)
etc...
```

This is an example which suggests both the benefits and the dangers. There is considerable benefit to having a copy of the data immediately available, but there also may be considerable expense involved in transporting it to all of the members of a list, if only a few will use the data anytime soon.

Alternatively, instead of replacing an external-body message with its real contents, it might make sense to transform it into a "multipart/alternative" message containing both the external body reference and the expanded version. This means that only the external body part can be forwarded if desired, and the recipient doesn't lose the information as to where the data was fetched from, if they want to fetch an up-to-date version in the future. Such information could be represented, in MIME, in the following form:

```
From: Nathaniel Borenstein <nsb@bellcore.com>
To: Ned Freed <ned@innosoft.com>
Subject: The latest MIME draft
Content-type: multipart/alternative; boundary=foo

--foo
Content-Type: message/external-body;
    name="BodyFormats.ps";
    site="thumper.bellcore.com";
    access-type=ANON-FTP;
    directory="pub";
    mode="image";
    expiration="Fri, 14 Jun 1991 19:13:14 -0400 (EDT)"

Content-type: application/postscript
--foo
Content-type: application/postscript

%!PS-Adobe-1.0
%%Creator: greenbush:nsb (Nathaniel Borenstein,MRE 2A-
274,4270,9938586,21462)
etc...
--foo--
```

Similarly for the case where a message is copied to a local FTP site, one could offer two external body parts as the alternatives, allowing the user agent to choose which FTP site is preferred.

Image and other Format Conversions

MIME currently defines two image formats, image/gif and image/jpeg. The former is much more convenient for many users, and can be displayed more quickly on many systems. The latter is a much more compact representation, and therefore places less stress on mail transport facilities.

Message transport services can optimize both transport bandwidth and user convenience by intelligent translation between these formats (and other formats that might be added later). When a message of type image/gif is submitted for long-haul delivery, it might reasonably be translated to image/jpeg. Conversely, when image/jpeg data is received for final delivery on a system with adequate storage resources, it might be translated to image/gif for the convenience of the recipient. Software to perform these translations is widely available. It should be noted, however, that performance of such conversions presumes support for the new format by the recipient.

Although MIME currently only defines one audio format, more are likely to be defined and registered with IANA in the future. In that case, similar format conversion facilities might be appropriate for audio.

If format conversion is done, it is **STRONGLY RECOMMENDED** that some kind of trace information (probably in the form of a Received header field) should be added to a message to document the conversion that has been performed.

Some people have expressed concerns, or even the opinion that conversions should never be done. To accommodate the desires of those who dislike the idea of automatic format conversion. For this reason, it is suggested that such transformations be generally restricted to gateways rather than general message transport services, and that services which perform such conversions should be sensitive to a header field that indicates that the sender does not wish to have any such conversions performed. A suggested value for this header field is:

Content-Conversion: prohibited

User agents that wish to explicitly indicate a willingness for such conversions to be performed may use:

Content-Conversion: permitted

However, this will be the default assumption of many gateways, so this header field is not strictly necessary. It also should be noted that such control of conversion would only be available to the sender, rather than to any of the recipients.

Robust Encoding of Data

In addition to all the reasons given above for possible transformation of body data, it will sometimes be the case that a gateway can tell that the body data, as given, will not robustly survive the next step of transport. For example, mail crossing an ASCII-to-EBCDIC gateway will lose information if certain characters are used. In such cases, a gateway can make the data more robust simply by applying one of the MIME Content-Transfer-Encoding algorithms (base64 or quoted-printable) to the body or body part. This will generally be a loss-less transformation, but care must be taken to ensure that the resulting message is MIME-conformant if the initial message was not. (For example, a MIME-Version header field may need to be added.)

User-oriented concerns

If a gateway is going to perform major transformations on a mail message, such as translating image formats or mapping between included data and external-reference data, it seems inevitable that there will be situations in which users will object to these transformations. This is, in large part, an implementation issue, but it seems advisable, wherever possible, to provide a mechanism whereby users can specify, to the transport system, whether or not they want such services performed automatically on their behalf. The use of the "Content-Conversion" header field, as mentioned above, is suggested for this purpose, since it at least provides some control by the sender, if not the recipient.

References

[RFC-1341] Borenstein, N., and N. Freed, "MIME (Multipurpose Internet Mail Extensions): Mechanisms for Specifying and Describing the Format of Internet Message Bodies", RFC 1341, Bellcore, June, 1992.

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Nathaniel S. Borenstein
MRE 2D-296, Bellcore
445 South St.
Morristown, NJ 07962-1910

Email: nsb@bellcore.com
Phone: +1 201 829 4270
Fax: +1 201 829 7019