

Internet Engineering Task Force (IETF)
Request for Comments: 6858
Updates: 3501
Category: Standards Track
ISSN: 2070-1721

A. Gulbrandsen
March 2013

Simplified POP and IMAP Downgrading for Internationalized Email

Abstract

This document specifies a method for IMAP and POP servers to serve internationalized messages to conventional clients. The specification is simple, easy to implement, and provides only rudimentary results.

Status of This Memo

This is an Internet Standards Track document.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6858>.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

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

1. Overview

A conventional IMAP or POP client may open a mailbox containing internationalized messages or may even attempt to read internationalized messages, for instance, when a user has both internationalized and conventional Mail User Agents (MUAs).

Some operations cannot be performed by conventional clients. Most importantly, an internationalized message usually contains at least one internationalized address, so address-based operations are rarely possible. This includes displaying the addresses, replying to messages, and the processing of most types of address-based signature or security.

However, the sender's name, message subject, body of text, and attachments can easily be displayed, so a helpful IMAP or POP server may prefer to display as much of the message as possible, rather than hide the message entirely.

This document specifies a way to present such messages to the client. It values simplicity of implementation over fidelity of representation, since implementing a high-fidelity downgrade algorithm such as the one specified in a companion document [RFC6857] is likely more work than implementing proper UTF-8 support for POP [RFC6856] and/or IMAP [RFC6855].

The server is assumed to be internationalized internally and to store messages that are internationalized messages natively. When it needs to present an internationalized message to a conventional client, the server synthesizes a conventional message containing most of the information and presents the "surrogate message".

This specification modifies the base IMAP specification [RFC3501] by relaxing a requirement that sizes be exact and adding a reporting requirement as discussed in Section 3 below.

2. Information Preserved and Lost

The surrogate message is intended to convey the most important information to the user. Where information is lost, the user should consider the message incomplete rather than modified.

The surrogate message is not intended to convey any information to the client software that would require or enable it to apply special handling to the message. Client authors who wish to handle internationalized messages are encouraged to implement POP [RFC6856] and/or IMAP [RFC6855] support for UTF-8.

Uppercase letters in examples represent non-ASCII characters. example.com is a plain domain; EXAMPLE.com represents a non-ASCII domain in the .com top-level domain.

2.1. Email Addresses

Each internationalized email address in the header fields listed below is replaced with an invalid email address whose display-name tells the user what happened.

The format of the display-name is explicitly unspecified. Anything that tells the user what happened is good. Anything that produces an email address that might belong to someone else is bad.

Given an internationalized address "Fred Foo <fred@EXAMPLE.com>", an implementation may choose to render it as one of these examples:

```
"fred@EXAMPLE.com" <invalid@internationalized-address.invalid>
Fred Foo <invalid@internationalized.invalid>
internationalized-address;;
fred;;
```

The .invalid top-level domain is reserved as a Top Level DNS Name [RFC2606]; therefore, the first two examples are syntactically valid, but they will never belong to anyone. Note that the display-name often needs encoding (see the Message Header Extensions document [RFC2047]).

The affected header fields are "Bcc:", "Cc:", "From:", "Reply-To:", "Resent-Bcc:", "Resent-Cc:", "Resent-From:", "Resent-Sender:", "Resent-To:", "Return-Path:", "Sender:", and "To:". Any addresses present in other header fields, such as "Received:", are not regarded as addresses by this specification.

2.2. MIME Parameters

Any MIME parameter [RFC2045] (whether in the message header or a body part header) that cannot be presented to the client exactly as it appears in the incoming message is silently excised.

Given a field such as

```
Content-Disposition: attachment; filename=FOO
```

the field is presented as

```
Content-Disposition: attachment
```

2.3. Subject Field

If the Subject field cannot be presented to the client exactly as it appears in the incoming message, the server presents a representation encoded as specified in RFC 2047.

2.4. Remaining Header Fields

Any header field that cannot be presented to the client, even with the modifications listed in Sections 2.1-2.3, is silently excised.

3. IMAP-Specific Details

IMAP allows clients to retrieve the message size without downloading the message, using RFC822.SIZE, BODY.SIZE[] and so on. The IMAP specification [RFC3501] requires that the returned size be exact.

This specification relaxes that requirement. When a conventional client requests size information for a message, the IMAP server is permitted to return size information for the internationalized message, even though the size of the surrogate message differs.

When an IMAP server performs downgrading as part of generating FETCH responses, it reports which messages were synthesized using a response code and attendant UID (Unique Identifier) set. This can be helpful to humans debugging the server and/or client.

```
C: a UID FETCH 1:* BODY.PEEK[HEADER.FIELDS(To From Cc)]
S: 1 FETCH (UID 65 [...])
S: 2 FETCH (UID 70 [...])
S: a OK [DOWNGRADED 70,105,108,109] Done
```

The message-set argument to DOWNGRADED contains UIDs.

Note that DOWNGRADED does not necessarily mention all the internationalized messages in the mailbox. In the example above, we know that UID 65 does not contain internationalized addresses in the "From:", "To:", and "Cc:" fields. It may, for example, contain an internationalized "Subject:".

4. POP-Specific Details

The number of lines specified in the TOP command [RFC1939] refers to the surrogate message. The message size reported by, for example, LIST may refer to either the internationalized or the surrogate message.

5. Security Considerations

If the internationalized message uses any sort of signature that covers header fields, the signature of the surrogate message almost certainly is invalid and may be invalid in other cases. This is a necessary limitation of displaying internationalized messages in legacy clients, since those clients do not support internationalized header fields. These cases are discussed in more detail in the POP or IMAP Downgrade document [RFC6857]. Even though invalid, these signatures should not be removed from the surrogate message, to preserve as much of the information as possible from the original message.

If any excised information is significant, then that information does not arrive at the recipient. Notably, the "Message-Id:", "In-Reference-To:", and "References:" fields may be excised, which might cause a lack of context when the recipient reads the message.

Some POP or IMAP clients, such as Fetchmail, download messages and delete the versions on the server. This may lead to permanent loss of information when the only remaining version of a message is the surrogate message.

Other clients cache messages for a very long time, even across client upgrades, such as the stock Android client. When such a client is internationalized, care must be taken so that it does not use an old surrogate message from its cache rather than retrieve the real message from the server.

6. IANA Considerations

IANA has added DOWNGRADED to the "IMAP Response Codes" registry.

7. References

7.1. Normative References

- [RFC1939] Myers, J. and M. Rose, "Post Office Protocol - Version 3", STD 53, RFC 1939, May 1996.
- [RFC2045] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.
- [RFC2047] Moore, K., "MIME (Multipurpose Internet Mail Extensions) Part Three: Message Header Extensions for Non-ASCII Text", RFC 2047, November 1996.

- [RFC2606] Eastlake, D., 3rd and A. Panitz, "Reserved Top Level DNS Names", BCP 32, RFC 2606, June 1999.
- [RFC3501] Crispin, M., "INTERNET MESSAGE ACCESS PROTOCOL - VERSION 4rev1", RFC 3501, March 2003.

7.2. Informative References

- [RFC1925] Callon, R., "The Twelve Networking Truths", RFC 1925, April 1 1996.
- [RFC6855] Resnick, P., Ed., Newman, C., Ed., and S. Shen, Ed., "IMAP Support for UTF-8", RFC 6855, March 2013.
- [RFC6856] Gellens, R., Newman, C., Yao, J., and K. Fujiwara, "Post Office Protocol Version 3 (POP3) Support for UTF-8", RFC 6856, March 2013.
- [RFC6857] Fujiwara, K., "Post-Delivery Message Downgrading for Internationalized Email Messages", RFC 6857, March 2013.

8. Acknowledgements

Claudio Allocchio, Ned Freed, Kazunori Fujiwara, Ted Hardie, John Klensin, Barry Leiba, John Levine, Alexey Melnikov, Chris Newman, and Joseph Yee. This specification was inspired by the principle stated in Rule 12 of "The Twelve Networking Truths" [RFC1925].

Author's Address

Arnt Gulbrandsen
Schweppermannstr. 8
D-81671 Muenchen
Germany

Fax: +49 89 4502 9758
EMail: arnt@gulbrandsen.priv.no

