

Network Working Group
Request for Comments: 5504
Category: Experimental

K. Fujiwara, Ed.
Y. Yoneya, Ed.
JPRS
March 2009

Downgrading Mechanism for Email Address Internationalization

Status of This Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

Traditional mail systems handle only ASCII characters in SMTP envelope and mail header fields. The Email Address Internationalization (UTF8SMTP) extension allows UTF-8 characters in SMTP envelope and mail header fields. To avoid rejecting internationalized email messages when a server in the delivery path does not support the UTF8SMTP extension, some sort of converting mechanism is required. This document describes a downgrading mechanism for Email Address Internationalization. Note that this is a way to downgrade, not tunnel. There is no associated up-conversion mechanism, although internationalized email clients might use original internationalized addresses or other data when displaying or replying to downgraded messages.

Table of Contents

1. Introduction	3
2. Terminology	4
3. New Header Fields Definition	5
3.1. Envelope Information Preservation Header Fields	5
3.2. Address Header Fields' Preservation Header Fields	6
3.3. Unknown Header Fields' Preservation Header Fields	6
4. SMTP Downgrading	7
4.1. Path Element Downgrading	7
4.2. ORCPT downgrading	8
5. Email Header Fields Downgrading	8
5.1. Downgrading Method for Each ABNF Element	8
5.1.1. RECEIVED Downgrading	9
5.1.2. UNSTRUCTURED Downgrading	9
5.1.3. WORD Downgrading	9
5.1.4. COMMENT Downgrading	9
5.1.5. MIME-VALUE Downgrading	9
5.1.6. DISPLAY-NAME Downgrading	9
5.1.7. MAILBOX Downgrading	9
5.1.8. ENCAPSULATION Downgrading	10
5.1.9. TYPED-ADDRESS Downgrading	10
5.2. Downgrading Method for Each Header Field	10
5.2.1. Address Header Fields That Contain <address>s	10
5.2.2. Address Header Fields with Typed Addresses	11
5.2.3. Downgrading Non-ASCII in Comments	11
5.2.4. Received Header Field	11
5.2.5. MIME Content Header Fields	12
5.2.6. Non-ASCII in <unstructured>	12
5.2.7. Non-ASCII in <phrase>	12
5.2.8. Other Header Fields	12
6. MIME Body-Part Header Field Downgrading	12
7. Security Considerations	13
8. Implementation Notes	14
8.1. RFC 2047 Encoding	14
8.2. Trivial Downgrading	15
8.3. 7bit Transport Consideration	15
9. IANA Considerations	16
10. Acknowledgements	18
11. References	18
11.1. Normative References	18
11.2. Informative References	19
Appendix A. Examples	20
A.1. Downgrading Example 1	20
A.2. Downgrading Example 2	22

1. Introduction

Traditional mail systems, which are defined by [RFC5321] and [RFC5322], allow ASCII characters in SMTP envelope and mail header field values. The UTF8SMTP extension ([RFC4952], [RFC5335], and [RFC5336]) allows UTF-8 characters in SMTP envelope and mail header field values.

If an envelope address or header field contains non-ASCII characters, the message cannot be delivered unless every system in the delivery path supports UTF8SMTP. This document describes a downgrading mechanism to avoid rejection of such messages when a server that does not support the UTF8SMTP extension is encountered. This downgrading mechanism converts envelope and mail header fields to an all-ASCII representation.

[RFC5335] allows UTF-8 characters to be used in mail header fields and MIME header fields. The downgrading mechanism specified here converts mail header fields and MIME header fields to ASCII.

This document does not change any protocols except by defining new header fields. It describes the conversion method from the internationalized email envelopes/messages that are defined in [RFC4952], [RFC5335], and [RFC5336] to the traditional email envelopes/messages defined in [RFC5321] and [RFC5322].

Section 3.2 of [RFC5336] defines when downgrading occurs. If the SMTP client has a UTF8SMTP envelope or an internationalized message and the SMTP server doesn't support the UTF8SMTP extension, then the SMTP client MUST NOT send a UTF8SMTP envelope or an internationalized message to the SMTP server. The section lists 4 choices in this case. The fourth choice is downgrading, as described here.

Downgrading may be implemented in Mail User Agents (MUAs), Mail Submission Agents (MSAs), and Mail Transport Agents (MTAs) that act as SMTP clients. It may also be implemented in Message Delivery Agents (MDAs), Post Office Protocol (POP) servers, and IMAP servers that store or offer UTF8SMTP envelopes or internationalized messages to non-UTF8SMTP-compliant systems, which include message stores.

This document tries to define the downgrading process clearly and it preserves the original internationalized email information as much as possible.

Downgrading in UTF8SMTP consists of the following four parts:

- o New header field definitions
- o SMTP downgrading
- o Email header field downgrading
- o MIME header field downgrading

In Section 3 of this document, many header fields starting with "Downgraded-" are introduced. They preserve the original envelope information and the original header fields.

SMTP downgrading is described in Section 4. It generates ASCII-only envelope information from a UTF8SMTP envelope.

Email header field downgrading is described in Section 5. It generates ASCII-only header fields.

MIME header fields are expanded in [RFC5335]. MIME header field downgrading is described in Section 6. It generates ASCII-only MIME header fields.

Displaying downgraded messages that originally contained internationalized email addresses or internationalized header fields is described in an another document ([DISPLAY]).

2. 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 [RFC2119].

All specialized terms used in this specification are defined in the Email Address Internationalization (EAI) overview [RFC4952], in the mail specifications [RFC5321] [RFC5322], or in the MIME documents [RFC2045] [RFC2047] [RFC2183] [RFC2231]. The terms "ASCII address", "internationalized email address", "non-ASCII address", "i18mail address", "UTF8SMTP", "message", and "mailing list" are used with the definitions from [RFC4952].

This document depends on [RFC5335], [RFC5336], and [RFC5337]. Key words used in those documents are used in this document, too.

The term "non-ASCII" refers to a UTF-8 string that contains at least one non-ASCII character.

A "UTF8SMTP envelope" has email originator/recipient addresses expanded by [RFC5336] and [RFC5337].

A "UTF8SMTP message" is an email message expanded by [RFC5335].

3. New Header Fields Definition

New header fields starting with "Downgraded-" are defined here to preserve those original envelope and mail header field values that contain UTF-8 characters. During downgrading, one new "Downgraded-" header field is added for each original envelope or mail header field that cannot be passed as-is to a server that does not support UTF8SMTP. The original envelope or mail header field is removed or rewritten. Only those envelope and mail header fields that contain non-ASCII characters are affected. The result of this process is a message that is compliant with existing email specifications [RFC5321] and [RFC5322]. The original internationalized information can be retrieved by examining the "Downgraded-" header fields that were added.

3.1. Envelope Information Preservation Header Fields

SMTP envelope downgraded information <downgraded-envelope-addr> consists of the original non-ASCII address and the downgraded all-ASCII address. The ABNF [RFC5234] syntax is as follows:

```
downgraded-envelope-addr = [FWS] "<" [ A-d-l ":" ] uMailbox
                           FWS "<" Mailbox ">" ">" [CFWS]
```

<uMailbox> is defined in [RFC5336]; <Mailbox> and <A-d-l> are defined in Section 4.1.2 of [RFC5321].

Two header fields, "Downgraded-Mail-From:" and "Downgraded-Rcpt-To:", are defined to preserve SMTP envelope downgraded information. The header field syntax is specified as follows:

```
fields                =/ downgradedmailfrom / downgradedrcptto

downgradedmailfrom =  "Downgraded-Mail-From:" unstructured CRLF

downgradedrcptto   =  "Downgraded-Rcpt-To:"   unstructured CRLF
```

The unstructured content is downgraded-envelope-addr and treated as if it were unstructured, with [RFC2047] encoding (and charset UTF-8) as needed.

3.2. Address Header Fields' Preservation Header Fields

The address header fields' preservation header fields are defined to preserve the original header field. Their value field holds the original header field value. The header field syntax is specified as follows:

```
fields                =/ known-downgraded-headers ":"
                        unstructured CRLF

known-downgraded-headers = "Downgraded-" original-headers

original-headers       = "From" / "Sender" /
                        "To" / "Cc" / "Bcc" /
                        "Reply-To" /
                        "Resent-From" / "Resent-Sender" /
                        "Resent-To" / "Resent-Cc" /
                        "Resent-Bcc" / "Resent-Reply-To" /
                        "Return-Path" /
                        "Disposition-Notification-To"
```

To preserve a header field in a "Downgraded-" header field:

1. Generate a new "Downgraded-" header field whose value is the original header field value.
2. Treat the generated header field content as if it were unstructured, and then apply [RFC2047] encoding with charset UTF-8 as necessary so that the result is ASCII.

3.3. Unknown Header Fields' Preservation Header Fields

The unknown header fields' preservation header fields are defined to encapsulate those original header fields that contain non-ASCII characters and are not otherwise provided for in this specification. The encapsulation header field name is the concatenation of "Downgraded-" and the original name. The value field holds the original header field value.

The header field syntax is specified as follows:

```
fields                =/ unknown-downgraded-headers ":" unstructured CRLF

unknown-downgraded-headers = "Downgraded-" original-header-field-name

original-header-field-name = field-name

field-name = 1*ftext
```

ftext = %d33-57 / ; Any character except
 %d59-126 ; controls, SP, and ":".

To encapsulate a header field in a "Downgraded-" header field:

1. Generate a new "Downgraded-" header field whose value is the original header field value.
2. Treat the generated header field content as if it were unstructured, and then apply [RFC2047] encoding with charset UTF-8 as necessary so the result is ASCII.
3. Remove the original header field.

4. SMTP Downgrading

The targets of downgrading elements in an SMTP envelope are below:

- o <reverse-path> of MAIL FROM command
- o <forward-path> of RCPT TO command
- o ORCPT parameter of RCPT TO command

<reverse-path> and <forward-path> are described in [RFC5321] and [RFC5336]. The ORCPT parameter is described in [RFC3461] and [RFC5337].

4.1. Path Element Downgrading

Downgrading the <path> of MAIL FROM and RCPT TO commands uses the ALT-ADDRESS parameter defined in [RFC5336]. An SMTP command is downgradable if the <path> contains a non-ASCII address and the command has an ALT-ADDRESS parameter that specifies an ASCII address. Since only non-ASCII addresses are downgradable, specifying an ALT-ADDRESS value for an all-ASCII address is invalid for use with this specification, and no interpretation is assigned to it. This restriction allows for future extension of the specification even though no such extensions are currently anticipated.

Note that even if no downgrading is performed on the envelope, message header fields and message body MIME header fields that contain non-ASCII characters MUST be downgraded. This is described in Sections 5 and 6.

When downgrading, replace each <path> that contains a non-ASCII mail address with its specified alternative ASCII address, and preserve the original information using "Downgraded-Mail-From" and

"Downgraded-Rcpt-To" header fields as defined in Section 3. Before replacing, decode the ALT-ADDRESS parameter value because it is encoded as xtext [RFC3461].

To avoid disclosing recipient addresses, the downgrading process MUST NOT add the "Downgraded-Rcpt-To:" header field if the SMTP downgrading targets multiple recipients. See Section 7 for more details.

As a result of the recipient address downgrading, the domain part of the recipient address prior to downgrading might be different from the domain part of the new recipient address. If the result of address resolution for the domain part of the new recipient address contains the server at the connection destination of the SMTP session for the recipient address prior to downgrading, the SMTP connection is valid for the new recipient address. Otherwise, the downgrading process MUST NOT send the downgraded message to the new recipient address via the connection and MUST try to send the downgraded message to the new recipient address.

4.2. ORCPT downgrading

The "RCPT TO" command can have an ORCPT parameter if the Delivery Status Notification (DSN) extension [RFC3461] is supported. If the ORCPT parameter contains a "utf-8" type address and the address contains raw non-ASCII characters, the address MUST be converted to utf-8-addr-xtext form. Those forms are described in [RFC5337] and clarified by successor documents such as [DSNBIS].

Before converting to utf-8-addr-xtext form, remove xtext encoding.

5. Email Header Fields Downgrading

This section defines the conversion method to ASCII for each header field that may contain non-ASCII characters.

[RFC5335] expands "Received:" header fields; [RFC5322] describes ABNF elements <mailbox>, <word>, <comment>, <unstructured>; [RFC2045] describes ABNF element <value>.

5.1. Downgrading Method for Each ABNF Element

Header field downgrading is defined below for each ABNF element. Downgrading an unknown header field is also defined as ENCAPSULATION downgrading. Converting the header field terminates when no non-ASCII characters remain in the header field.

5.1.1. RECEIVED Downgrading

If the header field name is "Received:" and the FOR clause contains a non-ASCII address, remove the FOR clause from the header field. Other parts (not counting <comment>s) should not contain non-ASCII values.

5.1.2. UNSTRUCTURED Downgrading

If the header field has an <unstructured> field that contains non-ASCII characters, apply [RFC2047] encoding with charset UTF-8.

5.1.3. WORD Downgrading

If the header field has any <word> fields that contain non-ASCII characters, apply [RFC2047] encoding with charset UTF-8.

5.1.4. COMMENT Downgrading

If the header field has any <comment> fields that contain non-ASCII characters, apply [RFC2047] encoding with charset UTF-8.

5.1.5. MIME-VALUE Downgrading

If the header field has any <value> elements defined by [RFC2045] and the elements contain non-ASCII characters, encode the <value> elements according to [RFC2231] with charset UTF-8 and leave the language information empty. If the <value> element is <quoted-string> and it contains <CFWS> outside the DQUOTE, remove the <CFWS> before this conversion.

5.1.6. DISPLAY-NAME Downgrading

If the header field has any <address> (<mailbox> or <group>) elements and they have <display-name> elements that contain non-ASCII characters, encode the <display-name> elements according to [RFC2047] with charset UTF-8. DISPLAY-NAME downgrading is the same algorithm as WORD downgrading.

5.1.7. MAILBOX Downgrading

The <mailbox> elements have no equivalent format for non-ASCII addresses. If the header field has any <mailbox> elements that contain non-ASCII characters, preserve the header field in the corresponding "Downgraded-" header field, which is defined in Section 3.2, and rewrite each <mailbox> element to ASCII-only format. The <mailbox> element that contains non-ASCII characters is one of three formats.

- o [Display-name] "<" Utf8-addr-spec 1*FCS "<" Addr-spec ">>"

Rewrite it as:

[Display-name] "<" Addr-spec ">"

- o [Display-name] "<" Utf8-addr-spec ">"
- o Utf8-addr-spec

Rewrite both as:

[Display-name] "Internationalized Address " Encoded-word
" Removed;";

where the <Encoded-word> is the original <Utf8-addr-spec>
encoded according to [RFC2047].

5.1.8. ENCAPSULATION Downgrading

If the header field contains non-ASCII characters and is such that no rule is given above, encapsulate it in a "Downgraded-" header field as described in Section 3.3 as a last resort.

Applying this procedure to "Received:" header field is prohibited.

5.1.9. TYPED-ADDRESS Downgrading

If the header field contains <utf-8-type-addr> and the <utf-8-type-addr> contains raw non-ASCII characters, it is in utf-8-address form. Convert it to utf-8-addr-xtext form as described in Section 4.2. COMMENT downgrading is also performed in this case. If the address type is unrecognized and the header field contains non-ASCII characters, then fall back to using ENCAPSULATION downgrading on the entire header field.

5.2. Downgrading Method for Each Header Field

Header fields are listed in [RFC4021]. This section describes the downgrading method for each header field.

If the whole mail header field does not contain non-ASCII characters, email header field downgrading is not required. Each header field's downgrading method is described below.

5.2.1. Address Header Fields That Contain <address>s

From:
Sender:
To:
Cc:
Bcc:

Reply-To:
Resent-From:
Resent-Sender:
Resent-To:
Resent-Cc:
Resent-Bcc:
Resent-Reply-To:
Return-Path:
Disposition-Notification-To:

If the header field contains <mailbox> elements that contain non-ASCII addresses, preserve the header field in a "Downgraded-" header field before the conversion. Then perform COMMENT downgrading, DISPLAY-NAME downgrading, and MAILBOX downgrading.

5.2.2. Address Header Fields with Typed Addresses

Original-Recipient:
Final-Recipient:

If the header field contains non-ASCII characters, perform TYPED-ADDRESS downgrading.

5.2.3. Downgrading Non-ASCII in Comments

Date:
Message-ID:
Resent-Message-ID:
In-Reply-To:
References:
Resent-Date:
Resent-Message-ID:
MIME-Version:
Content-ID:
Content-Transfer-Encoding:
Content-Language:
Accept-Language:
Auto-Submitted:

These header fields do not contain non-ASCII characters except in comments. If the header field contains UTF-8 characters in comments, perform COMMENT downgrading.

5.2.4. Received Header Field

Received:

Perform COMMENT downgrading and RECEIVED downgrading.

5.2.5. MIME Content Header Fields

Content-Type:
Content-Disposition:

Perform MIME-VALUE downgrading and COMMENT downgrading.

5.2.6. Non-ASCII in <unstructured>

Subject:
Comments:
Content-Description:

Perform UNSTRUCTURED downgrading.

5.2.7. Non-ASCII in <phrase>

Keywords:

Perform WORD downgrading.

5.2.8. Other Header Fields

For all other header fields that contain non-ASCII characters, are user-defined, and are missing from this document or future defined header fields, perform ENCAPSULATION downgrading.

If the software understands the header field's structure and a downgrading algorithm other than ENCAPSULATION is applicable, that software SHOULD use that algorithm; ENCAPSULATION downgrading is used as a last resort.

Mailing list header fields (those that start in "List-") are part of this category.

6. MIME Body-Part Header Field Downgrading

MIME body-part header fields may contain non-ASCII characters [RFC5335]. This section defines the conversion method to ASCII-only header fields for each MIME header field that contains non-ASCII characters. Parse the message body's MIME structure at all levels and check each MIME header field to see whether it contains non-ASCII characters. If the header field contains non-ASCII characters in the header field value, the header field is a target of the MIME body-part header field's downgrading. Each MIME header field's downgrading method is described below. COMMENT downgrading, MIME-VALUE downgrading, and UNSTRUCTURED downgrading are described in Section 5.

Content-ID:

The "Content-ID:" header field does not contain non-ASCII characters except in comments. If the header field contains UTF-8 characters in comments, perform COMMENT downgrading.

Content-Type:

Content-Disposition: Perform MIME-VALUE downgrading and COMMENT downgrading.

Content-Description: Perform UNSTRUCTURED downgrading.

7. Security Considerations

A downgraded message's header fields contain ASCII characters only. But they still contain MIME-encapsulated header fields that contain non-ASCII UTF-8 characters. Furthermore, the body part may contain UTF-8 characters. Implementations parsing Internet messages need to accept UTF-8 body parts and UTF-8 header fields that are MIME-encoded. Thus, this document inherits the security considerations of MIME-encoded header fields ([RFC2047] and [RFC3629]).

Rewriting header fields increases the opportunities for undetected spoofing by malicious senders. However, rewritten header fields are preserved into Downgraded-* header fields, and parsing Downgraded-* header fields enables the detection of spoofing caused by downgrading.

Addresses that do not appear in the message header fields may appear in the RCPT commands to an SMTP server for a number of reasons. Copying information from the envelope into the header fields risks inadvertent information disclosure (see [RFC5321] and Section 4 of this document). Mitigating inadvertent information disclosure is also discussed in these locations.

The techniques described here invalidate methods that depend on digital signatures over the envelope or any part of the message, which includes the top-level header fields and body-part header fields. Depending on the specific message being downgraded, the following techniques are likely to break: DomainKeys Identified Mail (DKIM), and possibly S/MIME and Pretty Good Privacy (PGP). The two obvious mitigations are to stick to 7-bit transport when using these techniques (as most/all of them presently require) or to make sure to have UTF8SMTP end-to-end when needed.

Many gateways and servers on the Internet will discard header fields with which they are not familiar. To the extent to which the downgrade procedures depend on new header fields (e.g.,

"Downgraded-") to avoid information loss, the risk of having those header fields dropped and subsequent implications must be identified. In particular, if the "Downgraded-" header fields are dropped, there is no possibility of reconstructing the original information at any point (before, during, or after delivery). Such gateways violate [RFC2979] and can be upgraded to correct the problem.

Even though the information is not lost, the original message cannot be perfectly reconstructed because some downgrading methods remove information (see Sections 5.1.1 and 5.1.5). Hence, downgrading is a one-way process.

While information in any email header field should usually be treated with some suspicion, current email systems commonly employ various mechanisms and protocols to make the information more trustworthy. Currently, information in the new Downgraded-* header fields is usually not inspected by these mechanisms, and may be even less trustworthy than the traditional header fields. Note that the Downgraded-* header fields could have been inserted with malicious intent (and with content unrelated to the traditional header fields).

If an internationalized MUA would simply try to "upgrade" the message for display purposes (that is, display the information in the Downgraded-* header fields instead of the traditional header fields), the effectiveness of the deployed mechanisms and protocols is likely to be reduced, and the user may be exposed to additional risks. More guidance on how to display downgraded messages is given in [DISPLAY].

Concerns about the trustworthiness of the Downgraded-* header fields are not limited to displaying and replying in MUAs, and should be carefully considered before using such header fields for other purposes as well.

See the "Security Considerations" section in [RFC4952] for more discussion.

8. Implementation Notes

8.1. RFC 2047 Encoding

While [RFC2047] has a specific algorithm to deal with whitespace in adjacent encoded words, there are a number of deployed implementations that fail to implement the algorithm correctly. As a result, whitespace behavior is somewhat unpredictable in practice when multiple encoded words are used. While RFC 5322 states that implementations SHOULD limit lines to not more than 78 characters, implementations MAY choose to allow overly long encoded words in

order to work around faulty [RFC2047] implementations. Implementations that choose to do so SHOULD have an optional mechanism to limit line length to 78 characters.

8.2. Trivial Downgrading

Downgrading is an alternative to avoid the rejection of messages that require UTF8SMTP support by a server that does not provide such support. Implementing the full specification of this document is desirable, but a partial implementation is also possible.

If a partial downgrading implementation confronts an unsupported downgrading target, the implementation MUST NOT send the message to a server that does not support UTF8SMTP. Instead, it MUST either reject the message or generate a notification of non-deliverability.

A partial downgrading, trivial downgrading, is discussed. It does not support non-ASCII addresses in SMTP envelope and address header fields, unknown header field downgrading, or the MIME body-part header field downgrading. It supports:

- o some simple header field downgrading: Subject
- o comments and display name downgrading: From, To, Cc
- o trace header field downgrading: Received

Otherwise, the downgrading fails.

Trivial downgrading targets mail messages that are generated by UTF8SMTP-aware MUAs and contain non-ASCII characters in comments, display names, and unstructured parts without using non-ASCII email addresses. These mail messages usually do not contain non-ASCII email addresses in the SMTP envelope and its header fields. But it is not deliverable via a UTF8SMTP-unaware SMTP server. Implementing full specification downgrading may be hard, but trivial downgrading saves mail messages without using non-ASCII addresses.

8.3. 7bit Transport Consideration

The SMTP client may encounter a SMTP server that does not support the 8BITMIME SMTP extension [RFC1652]. The server does not support "8bit" or "binary" data. Implementers need to consider converting "8bit" data to "base64" or "quoted-printable" encoded form and adjust the "Content-Transfer-Encoding" header field accordingly. If the body contains multiple MIME parts, this conversion MUST be performed for each MIME part.

9. IANA Considerations

IANA has registered the following header fields in the Permanent Message Header Field registry, in accordance with the procedures set out in [RFC3864].

Header field name: Downgraded-Mail-From
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Rcpt-To
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-From
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Sender
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-To
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Cc
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Bcc
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Reply-To
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Resent-From
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Resent-Sender
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Resent-To
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Resent-Cc
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Resent-Bcc
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Resent-Reply-To
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Return-Path
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Header field name: Downgraded-Disposition-Notification-To
Applicable protocol: mail
Status: experimental
Author/change controller: IETF
Specification document(s): This document (Section 3)

Furthermore, IANA is requested to refuse registration of all field names that start with "Downgraded-". For unknown header fields, use the downgrading method described in Section 3.3 to avoid conflicts with existing IETF activity (Email Address Internationalization).

10. Acknowledgements

Significant comments and suggestions were received from John Klensin, Harald Alvestrand, Chris Newman, Randall Gellens, Charles Lindsey, Marcos Sanz, Alexey Melnikov, Frank Ellermann, Edward Lewis, S. Moonesamy, and JET members.

11. References

11.1. Normative References

- [RFC1652] Klensin, J., Freed, N., Rose, M., Stefferud, E., and D. Crocker, "SMTP Service Extension for 8bit-MIMEtransport", RFC 1652, July 1994.
- [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.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2183] Troost, R., Dorner, S., and K. Moore, "Communicating Presentation Information in Internet Messages: The Content-Disposition Header Field", RFC 2183, August 1997.
- [RFC2231] Freed, N. and K. Moore, "MIME Parameter Value and Encoded Word Extensions: Character Sets, Languages, and Continuations", RFC 2231, November 1997.
- [RFC2979] Freed, N., "Behavior of and Requirements for Internet Firewalls", RFC 2979, October 2000.

- [RFC3461] Moore, K., "Simple Mail Transfer Protocol (SMTP) Service Extension for Delivery Status Notifications (DSNs)", RFC 3461, January 2003.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, September 2004.
- [RFC4021] Klyne, G. and J. Palme, "Registration of Mail and MIME Header Fields", RFC 4021, March 2005.
- [RFC4952] Klensin, J. and Y. Ko, "Overview and Framework for Internationalized Email", RFC 4952, July 2007.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.
- [RFC5335] Abel, Y., "Internationalized Email Headers", RFC 5335, September 2008.
- [RFC5336] Yao, J. and W. Mao, "SMTP Extension for Internationalized Email Addresses", RFC 5336, September 2008.
- [RFC5337] Newman, C. and A. Melnikov, "Internationalized Delivery Status and Disposition Notifications", RFC 5337, September 2008.

11.2. Informative References

- [DISPLAY] Fujiwara, K., "Displaying Downgraded Messages for Email Address Internationalization", Work in Progress, March 2009.
- [DSNBIS] Newman, C. and A. Melnikov, "Internationalized Delivery Status and Disposition Notifications", Work in Progress, December 2008.

Appendix A. Examples

A.1. Downgrading Example 1

This appendix shows an SMTP downgrading example. Consider a mail message where:

- o The sender address is "NON-ASCII-local@example.com", which is a non-ASCII address. Its ASCII alternative is "ASCII-local@example.com" and its display-name is "DISPLAY-local".
- o The "To:" address is "NON-ASCII-remotel@example.net", which is a non-ASCII address. Its ASCII alternative is "ASCII-remotel@example.net" and its display-name is "DISPLAY-remotel".
- o The "Cc:" address is a non-ASCII address, "NON-ASCII-remote2@example.org", without an alternative ASCII address. Its display-name is "DISPLAY-remote2".
- o Three display names contain non-ASCII characters.
- o The Subject header field is "NON-ASCII-SUBJECT", which contains non-ASCII characters.
- o Assume the "To:" recipient's MTA (example.net) does not support UTF8SMTP.
- o Assume the "Cc:" recipient's MTA (example.org) supports UTF8SMTP.

The first example SMTP envelope/message is shown in Figure 1. In this example, the "To:" recipient's session is the focus.

```
MAIL FROM: <NON-ASCII-local@example.com>
          ALT-ADDRESS=ASCII-local@example.com
RCPT TO:  <NON-ASCII-remotel@example.net>
          ALT-ADDRESS=ASCII-remotel@example.net
RCPT TO:  <NON-ASCII-remote2@example.org>
-----
Message-Id: MESSAGE_ID
Mime-Version: 1.0
Content-Type: text/plain; charset="UTF-8"
Content-Transfer-Encoding: 8bit
Subject: NON-ASCII-SUBJECT
From: DISPLAY-local <NON-ASCII-local@example.com
      <ASCII-local@example.com>>
To: DISPLAY-remotel <NON-ASCII-remotel@example.net
    <ASCII-remotel@example.net>>
Cc: DISPLAY-remote2 <NON-ASCII-remote2@example.org>
Date: DATE
```

MAIL_BODY

Figure 1: Original envelope/message (example 1)

In this example, there are two SMTP recipients; one is "To:", the other is "Cc:". The SMTP downgrading uses To: session downgrading. Figure 2 shows an SMTP downgraded example.

```
MAIL FROM: <ASCII-local@example.com>
RCPT TO:  <ASCII-remotel@example.net>
-----
Downgraded-Mail-From: <NON-ASCII-local@example.com <ASCII-local@example.com>>   Downgr
aded-Rcpt-To: <NON-ASCII-remotel@example.net <ASCII-remotel@example.net>>   Message-Id: M
ESSAGE_ID
Mime-Version: 1.0
Content-Type: text/plain; charset="UTF-8"
Content-Transfer-Encoding: 8bit
Subject: NON-ASCII-SUBJECT
From: DISPLAY-local <NON-ASCII-local@example.com
      <ASCII-local@example.com>>
To: DISPLAY-remotel <NON-ASCII-remotel@example.net
    <ASCII-remotel@example.net>>
Cc: DISPLAY-remote2 <NON-ASCII-remote2@example.org>
Date: DATE
```

MAIL_BODY

Figure 2: SMTP downgraded envelope/message (example 1)

After SMTP downgrading, header field downgrading is performed. The final downgraded message is shown in Figure 3. A Return-Path header field will be added by the final destination MTA.

```
Return-Path: <ASCII-local@example.com>
Downgraded-Mail-From: <NON-ASCII-local@example.com <ASCII-local@example.com>>
Downgraded-Rcpt-To: <NON-ASCII-remote1@example.net <ASCII-remote1@example.net>>
Message-Id: MESSAGE_ID
Mime-Version: 1.0
Content-Type: text/plain; charset="UTF-8"
Content-Transfer-Encoding: 8bit
Subject: NON-ASCII-SUBJECT
From: DISPLAY-local <ASCII-local@example.com>
Downgraded-From: DISPLAY-local <NON-ASCII-local@example.com <ASCII-local@example.com>>
To: DISPLAY-remote1 <ASCII-remote1@example.net>
Downgraded-To: DISPLAY-remote1 <NON-ASCII-remote1@example.net <ASCII-remote1@example.net>
>
Cc: DISPLAY-remote2 Internationalized address
  NON-ASCII-remote2@example.org removed;;
Downgraded-Cc: DISPLAY-remote2 <NON-ASCII-remote2@example.org>
Date: DATE

MAIL_BODY
```

Figure 3: Downgraded message (example 1)

A.2. Downgrading Example 2

In many cases, the sender wants to use a non-ASCII address and the recipient is a traditional mail user. The SMTP server handing mail for the recipient and/or the recipient's MUA does not support UTF8SMTP extension. Consider a mail message where:

- o The sender address is "NON-ASCII-local@example.com", which is a non-ASCII address. Its ASCII alternative is "ASCII-local@example.com". It has a display-name "DISPLAY-local", which contains non-ASCII characters.
- o The "To:" address is "ASCII-remote1@example.net", which is ASCII-only. It has a display-name, "DISPLAY-remote1", which contains non-ASCII characters.
- o The "Subject:" header field is "NON-ASCII-SUBJECT", which contains non-ASCII characters.

The second example envelope/message is shown in Figure 4.

```
MAIL From: <NON-ASCII-local@example.com>
          ALT-ADDRESS=ASCII-local@example.com
RCPT TO: <ASCII-remotel@example.net>
-----
Message-Id: MESSAGE_ID
Mime-Version: 1.0
Content-Type: text/plain; charset="UTF-8"
Content-Transfer-Encoding: 8bit
Subject: NON-ASCII-SUBJECT
From: DISPLAY-local <NON-ASCII-local@example.com
     <ASCII-local@example.com>>
To: DISPLAY-remotel <ASCII-remotel@example.net>
Date: DATE

MAIL_BODY
```

Figure 4: Original message (example 2)

In this example, SMTP session is downgradable. Figure 5 shows an SMTP downgraded envelope/message.

```
MAIL From: <ASCII-local@example.com>
RCPT TO: <ASCII-remotel@example.net>
-----
Downgraded-Mail-From: <NON-ASCII-local@example.com
om>>?=?=UTF8?Q?<ASCII-local@example.c
Message-Id: MESSAGE_ID
Mime-Version: 1.0
Content-Type: text/plain; charset="UTF-8"
Content-Transfer-Encoding: 8bit
Subject: NON-ASCII-SUBJECT
From: DISPLAY-local <NON-ASCII-local@example.com
     <ASCII-local@example.com>>
To: DISPLAY-remotel <ASCII-remotel@example.net>
Date: DATE

MAIL_BODY
```

Figure 5: SMTP downgraded envelope/message (example 2)

After SMTP downgrading, header field downgrading is performed. The downgraded example is shown in Figure 6.

```
Return-Path: <ASCII-local@example.com>
Downgraded-Mail-From: <NON-ASCII-local@example.com =?UTF8?Q?<ASCII-local@example.com>>?=
Message-Id: MESSAGE_ID
Mime-Version: 1.0
Content-Type: text/plain; charset="UTF-8"
Content-Transfer-Encoding: 8bit
Subject: NON-ASCII-SUBJECT
Downgraded-From: DISPLAY-local <NON-ASCII-local@example.com <ASCII-local@example.com>>
From: DISPLAY-local <ASCII-local@example.com>
To: DISPLAY-remotel <ASCII-remotel@example.net>
Date: DATE
```

MAIL_BODY

Figure 6: Downgraded message (example 2)

Authors' Addresses

Kazunori Fujiwara (editor)
Japan Registry Services Co., Ltd.
Chiyoda First Bldg. East 13F, 3-8-1 Nishi-Kanda
Chiyoda-ku, Tokyo 101-0065
Japan

Phone: +81 3 5215 8451
EMail: fujiwara@jprs.co.jp

Yoshiro Yoneya (editor)
Japan Registry Services Co., Ltd.
Chiyoda First Bldg. East 13F, 3-8-1 Nishi-Kanda
Chiyoda-ku, Tokyo 101-0065
Japan

Phone: +81 3 5215 8451
EMail: yone@jprs.co.jp

