

Network Working Group
Internet-Draft
Updates: [3461](#), 3464, 3798
(if approved)
Intended status: Experimental
Expires: January 2, 2008

C. Newman
Sun Microsystems
A. Melnikov, Ed.
Isode Ltd
July 1, 2007

International Delivery and Disposition Notifications
draft-ietf-eai-dsn-02.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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 and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 2, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

Delivery status notifications (DSNs) are critical to the correct operation of an email system. However, the existing draft standard is presently limited to US-ASCII text in the machine readable portions of the protocol. This specification adds a new address type for international email addresses so an original recipient address with non-US-ASCII characters can be correctly preserved even after

downgrading. This also provides updated content return media types for delivery status notifications and message disposition notifications to support use of the new address type.

This document experimentally extends [RFC 3461](#), [RFC 3464](#) and [RFC 3798](#).

Table of Contents

1.	Introduction	3
2.	Conventions Used in this Document	3
3.	UTF-8 Address Type	3
4.	UTF-8 Delivery Status Notifications	6
5.	UTF-8 Message Disposition Notifications	8
6.	IANA Considerations	9
6.1.	UTF-8 Mail Address Type Registration	9
6.2.	Update to 'smtp' Diagnostic Type Registration	10
6.3.	message/utf-8-headers	10
6.4.	message/utf-8-delivery-status	11
6.5.	message/utf-8-disposition-notification	13
7.	Security Considerations	14
8.	References	15
8.1.	Normative References	15
8.2.	Informative References	16
Appendix A.	Acknowledgements	16
Appendix B.	Open Issues	16
Appendix C.	Changes from -01	16
Appendix D.	Changes from -00	17
	Authors' Addresses	17
	Intellectual Property and Copyright Statements	18

1. Introduction

When an email message is transmitted using the UTF8SMTP [[I-D.ietf-eai-smtpext](#)] extension and Internationalized Email Headers [[I-D.ietf-eai-utf8headers](#)], it is sometimes necessary to return that message or generate a Message Disposition Notification [[RFC3798](#)] (MDN). As a message sent to multiple recipients can generate a status and disposition notification for each recipient, it is helpful if a client can correlate these returns based on the recipient address it provided, thus preservation of the original recipient is important. This specification describes how to preserve the original recipient and updates the MDN and DSN formats to support the new address types.

2. Conventions Used in this Document

The key words "MUST", "MUST NOT", "SHOULD", "SHOULD NOT", and "MAY" in this document are to be interpreted as defined in "Key words for use in RFCs to Indicate Requirement Levels" [[RFC2119](#)].

The formal syntax use the Augmented Backus-Naur Form (ABNF) [[RFC4234](#)] notation including the core rules defined in [Appendix B of RFC 4234](#) and the rules in [section 4 of RFC 3629](#).

3. UTF-8 Address Type

An Extensible Message Format for Delivery Status Notifications [[RFC3464](#)] defines the concept of an address type. The address format introduced in Internationalized Email Headers [[I-D.ietf-eai-utf8headers](#)] is a new address type. The syntax for the new address type in the context of status notifications follows:

An SMTP [[RFC2821](#)] server which advertises both the UTF8SMTP extension [[I-D.ietf-eai-smtpext](#)] and the DSN extension [[RFC3461](#)] MUST accept a utf-8 address type in the ORCPT parameter including 8-bit UTF-8 characters. This address type also includes a 7-bit encoding suitable for use in a message/delivery-status body part or an ORCPT parameter sent to an SMTP server which does not advertise UTF8SMTP.

This address type has 3 forms: utf-8-addr-xtext, utf-8-addr-unitext and utf-8-address. The first 2 forms are 7-bit safe.

The utf-8-address form is only suitable for use in newly defined protocols capable of native representation of 8-bit characters. I.e. the utf-8-address form MUST NOT be used in the ORCPT parameter when the SMTP server doesn't advertise support for UTF8SMTP or the SMTP

server supports UTF8SMTP, but the address contains US-ASCII characters not permitted in the ORCPT parameter (e.g. the ORCPT parameter forbids SP and =); or in a 7-bit transport environment including a message/delivery-status Original-Recipient or Final-Recipient field. The utf-8-addr-xtext form (see below) MUST be used instead in the former case, the utf-8-addr-unitext form MUST be used in the latter case. The utf-8-address form MAY be used in the ORCPT parameter when the SMTP server also advertises support for UTF8SMTP and the address doesn't contain any US-ASCII characters not permitted in the ORCPT parameter. It SHOULD be used in a message/utf-8-delivery-status Original-Recipient or Final-Recipient DSN field; or in an Original-Recipient header field [[RFC3798](#)] if the message is a UTF-8 header message.

In addition, the utf-8-addr-unitext form can be used anywhere where the utf-8-address form is allowed.

When using in the ORCPT parameter, the utf-8 address type requires that US-ASCII CTLs, SP, \, + and = be encoded using xtext encoding as described in [[RFC3461](#)]. This is described by the utf-8-addr-xtext form in the ABNF below. Unicode characters MAY be included in a utf-8 address type using a "\x{HEXPOINT}" syntax, where HEXPOINT is 2 to 6 hexadecimal digits. When sending data to a UTF8SMTP capable server, native UTF-8 characters SHOULD be used instead of the QMIDCHAR and QHIGHCHAR encodings described below. When sending data to an SMTP server which does not advertise UTF8SMTP, then the QMIDCHAR and QHIGHCHAR encodings MUST be used instead of UTF-8.

When the ORCPT parameter is placed in a message/utf-8-delivery-status Original-Recipient field, the utf-8-addr-xtext form of the utf-8 address type SHOULD be converted to the 'utf-8-address' form (see the ABNF below) by removing all xtext encoding first (which will result in the 'utf-8-addr-unitext' form), followed by removal of the 'unitext' encoding. However, if an address is labeled with the utf-8 address type but does not conform to utf-8 syntax, then it MUST be copied into the message/utf-8-delivery-status field without alteration.

The ability to encode characters with the EmbeddedUnicodeChar encodings should be viewed as a transitional mechanism. It is hoped that as systems lacking support for UTF8SMTP become less common over time, these encodings can eventually be phased out.


```
utf-8-type-addr      = "utf-8;" utf-8-enc-addr
```

```
utf-8-address        = uMailbox [ *WSP "<" Mailbox ">" ]
; 'uMailbox' is defined in [I-D.ietf-eai-smtpext].
; 'Mailbox' is defined in [RFC2821].
```

```
utf-8-enc-addr       = utf-8-addr-xtext /
                      utf-8-addr-unitext /
                      utf-8-address
```

```
///Add comment about which where each type is used
```

```
utf-8-addr-xtext     = xtext
; xtext is defined in [RFC3461].
; When xtext encoding is removed,
; the syntax MUST conform to
; 'utf-8-addr-unitext'.
```

```
utf-8-addr-unitext   = 1*(QUCHAR / EmbeddedUnicodeChar)
; MUST follow 'utf-8-address' ABNF when
; dequoted
```

```
QUCHAR               = %x21-2a / %x2c-3c / %x3e-5b / %x5d-7e /
                      UTF8-2 / UTF8-3 / UTF8-4
; Printable except CTLs, SP, '\', '+' and '='
```

```
EmbeddedUnicodeChar = %x5C.78 "{" HEXPOINT "}"
; starts with "\x"
```

```
HEXPOINT = "5C" / (HEXDIG8 HEXDIG) /      ; 2 digit forms
( NZHEXDIG 2(HEXDIG) ) /                  ; 3 digit forms
( NZDHEXDIG 3(HEXDIG) ) /
( "D" %x30-37 2(HEXDIG) ) /
; 4 digit forms excluding surrogate
( NZHEXDIG 4(HEXDIG) ) /                  ; 5 digit forms
( "10" 4*HEXDIG )                        ; 6 digit forms
; represents either "\" or a Unicode code point outside the
; US-ASCII repertoire
```

```
HEXDIG8              = %x38-39 / "A" / "B" / "C" / "D" / "E" / "F"
; HEXDIG excluding 0-7
```

```
NZHEXDIG             = %x31-39 / "A" / "B" / "C" / "D" / "E" / "F"
; HEXDIG excluding "0"
```

```
NZDHEXDIG            = %x31-39 / "A" / "B" / "C" / "E" / "F"
; HEXDIG excluding "0" and "D"
```


4. UTF-8 Delivery Status Notifications

A traditional delivery status notification [[RFC3464](#)] comes in a three-part multipart/report [[RFC3462](#)] container, where the first part is human readable text describing the error, the second part is a 7-bit-only message/delivery-status and the optional third part is used for content (message/rfc822) or header (text/rfc822-headers) return. An SMTP server which advertises both UTF8SMTP and DSN SHOULD return an undeliverable UTF8SMTP message without downgrading it (assuming the return SMTP server supports UTF8SMTP). As the present DSN format does not permit this, three new media types are needed.

The first type, message/utf-8-delivery-status has the syntax of message/delivery-status with three modifications. First, the charset for message/utf-8-delivery-status is UTF-8 and thus any field MAY contain UTF-8 characters when appropriate (see the ABNF below). In particular, the Diagnostic-Code field MAY contain UTF-8 as described in UTF8SMTP [[I-D.ietf-eai-smtpext](#)]; the Diagnostic-Code field SHOULD be in i-default language [[LANGTAGS](#)]. Second, systems generating a message/utf-8-delivery-status body part SHOULD use the utf-8-address form of the utf-8 address type for all addresses containing characters outside the US-ASCII repertoire. These systems SHOULD up-convert the utf-8-addr-xttext or the utf-8-addr-unitext form of a utf-8 address type in the ORCPT parameter to the utf-8-address form of a utf-8 address type in the Original-Recipient field. Third, a new optional field called Localized-Diagnostic is added. Each instance includes a language tag [[LANGTAGS](#)] and contains text in the specified language. This is equivalent to the text part of the Diagnostic-Code field. All instances of Localized-Diagnostic MUST use different language tags. The ABNF for message/utf-8-delivery-status is specified below:


```
utf-8-delivery-status-content = per-message-fields
                                1*( CRLF utf-8-per-recipient-fields )
    ; "per-message-fields" remains unchanged from the definition
    ; in RFC 3464, except for the "extension-field"
    ; which is updated below.

utf-8-per-recipient-fields =
    [ original-recipient-field CRLF ]
    final-recipient-field CRLF
    action-field CRLF
    status-field CRLF
    [ remote-mta-field CRLF ]
    [ diagnostic-code-field CRLF
      *(localized-diagnostic-text-field CRLF) ]
    [ last-attempt-date-field CRLF ]
    [ will-retry-until-field CRLF ]
    *( extension-field CRLF )
    ; All fields except for "original-recipient-field",
    ; "final-recipient-field", "diagnostic-code-field"
    ; and "extension-field" remain unchanged from
    ; the definition in RFC 3464.

generic-address =/ utf-8-enc-addr
    ; Only allowed with the "utf-8" address-type.
    ;
    ; This indirectly updates "original-recipient-field"
    ; and "final-recipient-field"

diagnostic-code-field =
    "Diagnostic-Code" ":" diagnostic-type ";" *text

localized-diagnostic-text-field =
    "Localized-Diagnostic" ":" Language-Tag ";" *utf8-text
    ; "Language-Tag" is a language tag as defined in [LANGTAGS].

extension-field =/ extension-field-name ":" *utf8-text

utf8-text = %d1-9 /           ; Any Unicode character except for NUL,
                %d11 /        ; CR and LF, encoded in UTF-8
                %d12 /
                %d14-127 /
                UTFMB

UTFMB      = UTF2 / UTF3 / UTF4
```

The second type, used for content return, is message/utf-8 which is similar to message/rfc822, except it contains a message with UTF-8 headers. This media type is described in [[I-D.ietf-eai-utf8headers](#)].

The third type, used for header return, is message/utf-8-headers and contains only the UTF-8 headers of a message (all lines prior to the first blank line in a UTF8SMTP message). Unlike message/utf-8, this body part provides no difficulties for present infrastructure.

All three new types will typically use the "8bit" Content-Transfer-Encoding (in the event all content is 7-bit, the equivalent traditional types for delivery status notifications are advised for greater backwards compatibility). While MIME [RFC2046] advises against the use of 8-bit in new message subtypes intended for the email infrastructure, that advice does not apply to these new types which are intended primarily for use by newer systems with full support for 8-bit MIME and UTF-8 headers.

5. UTF-8 Message Disposition Notifications

Message Disposition Notifications [RFC3798] have a similar design and structure to DSNs. As a result, they use the same basic return format. When generating a MDN for a UTF-8 header message, content or header return is the same as for DSNs. The second part of the multipart/report uses a new media type, message/utf-8-disposition-notification, which has the syntax of message/disposition-notification with two modifications. First, the charset for message/utf-8-disposition-notification is UTF-8 and thus any field MAY contain UTF-8 characters when appropriate (see the ABNF below). (In particular, the failure-field, the error-field and the warning-field MAY contain UTF-8. These fields SHOULD be in i-default language [LANGTAGS].) Second, systems generating a message/utf-8-disposition-notification body part (typically a mail user agent) SHOULD use the utf-8 address type for all addresses containing characters outside the US-ASCII repertoire.

The MDN specification also defines the Original-Recipient header field which is added with a copy of the contents of ORCPT at delivery time. When generating an Original-Recipient header field, a delivery agent writing a UTF-8 header message in native format SHOULD convert the utf-8-addr-xttext or the utf-8-addr-unitext form of a utf-8 address type in the ORCPT parameter to the corresponding utf-8-address form.

The MDN specification also defines the Disposition-Notification-To header which is an address header and thus follows the same 8-bit rules as other address headers such as "From" and "To" when used in a UTF-8 header message.

; ABNF for "original-recipient-header", "original-recipient-field"
; and "final-recipient-field" from [RFC 3798](#) is implicitly updated
; as they use the updated "generic-address" as defined in
; [section 4](#) of this document.

failure-field = "Failure" ":" *utf8-text
; "utf8-text" is defined in [section 4](#) of this document.

error-field = "Error" ":" *utf8-text
; "utf8-text" is defined in [section 4](#) of this document.

warning-field = "Warning" ":" *utf8-text
; "utf8-text" is defined in [section 4](#) of this document.

6. IANA Considerations

This specification does not create any new IANA registries. However the following items are registered as a result of this document:

6.1. UTF-8 Mail Address Type Registration

The mail address type registry was created by [RFC 3464](#). The registration template response follows:

(a) The proposed address-type name.

UTF-8

(b) The syntax for mailbox addresses of this type, specified using BNF, regular expressions, ASN.1, or other non-ambiguous language.

See [Section 3](#).

(c) If addresses of this type are not composed entirely of graphic characters from the US-ASCII repertoire, a specification for how they are to be encoded as graphic US-ASCII characters in a DSN Original-Recipient or Final-Recipient DSN field.

This address type has 3 forms (as defined in [Section 3](#)): utf-8-addr-text, utf-8-addr-unitext and utf-8-address. The first 2 forms are 7-bit safe.

The utf-8-address form MUST NOT be used in the ORCPT parameter when the SMTP server doesn't advertise support for UTF8SMTP or the SMTP server supports UTF8SMTP, but the address contains US-ASCII characters not permitted in the ORCPT parameter (e.g. the ORCPT parameter forbids SP and =); or in a 7-bit transport environment

including a message/delivery-status Original-Recipient or Final-Recipient field. The utf-8-addr-xttext form MUST be used instead in the former case, the utf-8-addr-unixtext form MUST be used in the latter case. The utf-8-address form MAY be used in the ORCPT parameter when the SMTP server also advertises support for UTF8SMTP and the address doesn't contain any US-ASCII characters not permitted in the ORCPT parameter; in a message/utf-8-delivery-status Original-Recipient or Final-Recipient DSN field; or an Original-Recipient header field [[RFC3798](#)] if the message is a UTF-8 header message.

In addition, the utf-8-addr-unixtext form can be used anywhere where the utf-8-address form is allowed.

[6.2.](#) Update to 'smtp' Diagnostic Type Registration

The mail diagnostic type registry was created by [RFC 3464](#). The registration for the 'smtp' diagnostic type should be updated to reference RFC XXXX in addition to [RFC 3464](#).

When the 'smtp' diagnostic type is used in the context of a message/delivery-status body part, it remains as presently defined. When the 'smtp' diagnostic type is used in the context of a message/utf-8-delivery-status body part, the codes remain the same, but the text portion MAY contain UTF-8 characters.

[6.3.](#) message/utf-8-headers

Type name: message

Subtype name: utf-8-headers

Required parameters: none

Optional parameters: none

Encoding considerations: This media type contains Internationalized Email Headers [[I-D.ietf-eai-utf8headers](#)] with no message body. Whenever possible, the 8-bit content transfer encoding SHOULD be used. When this media type passes through a 7-bit-only SMTP infrastructure it MAY be encoded with the base64 or quoted-printable content transfer encoding.

Security considerations: See [Section 7](#)

Interoperability considerations: It is important this media type is not converted to a charset other than UTF-8. As a result, implementations MUST NOT include a charset parameter with this media type. Although it might be possible to downconvert this media type to the text/rfc822-header media type, such conversion is discouraged as it loses information.

Published specification: RFC XXXX

Applications that use this media type: UTF8SMTP servers and email clients that support multipart/report generation or parsing.

Additional information:

Magic number(s): none

File extension(s): In the event this is saved to a file, the extension ".u8hdr" is suggested.

Macintosh file type code(s): The 'TEXT' type code is suggested as files of this type are typically used for diagnostic purposes and suitable for analysis in a UTF-8 aware text editor. A uniform type identifier (UTI) of "public.utf8-email-message-header" is suggested. This type conforms to "public.utf8-plain-text" and "public.plain-text".

Person & email address to contact for further information: See the Author's address section of this document.

Intended usage: COMMON

Restrictions on usage: This media type contains textual data in the UTF-8 charset. It typically contains octets with the 8th bit set. As a result a transfer encoding is required when a 7-bit transport is used.

Author: See Author's Address section of this document.

Change controller: IETF Standards Process

6.4. message/utf-8-delivery-status

Type name: message

Subtype name: utf-8-delivery-status

Required parameters: none

Optional parameters: none

Encoding considerations: This media type contains delivery status notification attributes in the UTF-8 charset. The 8-bit content transfer encoding **MUST** be used with this content-type, unless it is sent over a 7-bit transport environment in which case quoted-printable or base64 may be necessary.

Security considerations: See [Section 7](#)

Interoperability considerations: This media type provides functionality similar to the message/delivery-status content type for email message return information. Clients of the previous format will need to be upgraded to interpret the new format, however the new media type makes it simple to identify the difference.

Published specification: RFC XXXX

Applications that use this media type: SMTP servers and email clients that support delivery status notification generation or parsing.

Additional information:

Magic number(s): none

File extension(s): The extension ".u8dsn" is suggested.

Macintosh file type code(s): A uniform type identifier (UTI) of "public.utf8-email-message-delivery-status" is suggested. This type conforms to "public.utf8-plain-text".

Person & email address to contact for further information: See the Author's address section of this document.

Intended usage: COMMON

Restrictions on usage: This is expected to be the second part of a multipart/report.

Author: See Author's Address section of this document.

Change controller: IETF Standards Process

6.5. message/utf-8-disposition-notification

Type name: message

Subtype name: utf-8-disposition-notification

Required parameters: none

Optional parameters: none

Encoding considerations: This media type contains disposition notification attributes in the UTF-8 charset. The 8-bit content transfer encoding **MUST** be used with this content-type, unless it is sent over a 7-bit transport environment in which case quoted-printable or base64 may be necessary.

Security considerations: See [Section 7](#)

Interoperability considerations: This media type provides functionality similar to the message/disposition-notification content type for email message disposition information. Clients of the previous format will need to be upgraded to interpret the new format, however the new media type makes it simple to identify the difference.

Published specification: RFC XXXX

Applications that use this media type: Email clients or servers that support message disposition notification generation or parsing.

Additional information:

Magic number(s): none

File extension(s): The extension ".u8mdn" is suggested.

Macintosh file type code(s): A uniform type identifier (UTI) of "public.utf8-email-message-disposition-notification" is suggested. This type conforms to "public.utf8-plain-text".

Person & email address to contact for further information: See the Author's address section of this document.

Intended usage: COMMON

Restrictions on usage: This is expected to be the second part of a multipart/report.

Author: See Author's Address section of this document.

Change controller: IETF Standards Process

7. Security Considerations

Automated use of report types without authentication presents several security issues. Forging negative reports presents the opportunity for denial-of-service attacks when the reports are used for automated maintenance of directories or mailing lists. Forging positive reports may cause the sender to incorrectly believe a message was delivered when it was not.

Malicious users can generate report structures designed to trigger coding flaws in report parsers. Report parsers need to use secure coding techniques to avoid the risk of buffer overflow or denial-of-service attacks against parser coding mistakes. Code reviews of such parsers are also recommended.

Malicious users of the email system regularly send messages with forged envelope return paths and these messages trigger delivery status reports that result in a large amount of unwanted traffic on the Internet. Many users choose to ignore delivery status notifications because they are usually the result of "blowback" from forged messages and thus never notice when messages they sent go undelivered. As a result, support for correlation of delivery status and message disposition notification messages with sent-messages has become a critical feature of mail clients and possibly mail stores if the email infrastructure is to remain reliable. In the short term, simply correlating message-IDs may be sufficient to distinguish true status notifications from those resulting from forged originator addresses. But in the longer term, including cryptographic signature material that can securely associate the status notification with the original message is advisable.

As this specification permits UTF-8 in additional fields, the security considerations of UTF-8 [[RFC3629](#)] apply.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2821] Klensin, J., "Simple Mail Transfer Protocol", [RFC 2821](#), April 2001.
- [RFC3461] Moore, K., "Simple Mail Transfer Protocol (SMTP) Service Extension for Delivery Status Notifications (DSNs)", [RFC 3461](#), January 2003.
- [RFC3462] Vaudreuil, G., "The Multipart/Report Content Type for the Reporting of Mail System Administrative Messages", [RFC 3462](#), January 2003.
- [RFC3464] Moore, K. and G. Vaudreuil, "An Extensible Message Format for Delivery Status Notifications", [RFC 3464](#), January 2003.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](#), November 2003.
- [RFC3798] Hansen, T. and G. Vaudreuil, "Message Disposition Notification", [RFC 3798](#), May 2004.
- [RFC4234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", [RFC 4234](#), October 2005.
- [I-D.ietf-eai-utf8headers]
Yang, A., "Internationalized Email Headers",
[draft-ietf-eai-utf8headers-05](#) (work in progress),
April 2007.
- [I-D.ietf-eai-smtpext]
Yao, J. and W. Mao, "SMTP extension for internationalized email address", [draft-ietf-eai-smtpext-05](#) (work in progress), April 2007.
- [I-D.ietf-eai-downgrade]
Yoneya, Y. and K. Fujiwara, "Downgrading mechanism for Email Address Internationalization (EAI)",
[draft-ietf-eai-downgrade-03](#) (work in progress), Mar 2007.
- [LANGTAGS]
Phillips, A. and M. Davis, "Tags for Identifying Languages", [RFC 4646](#), September 2006.

8.2. Informative References

- [RFC2045] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", [RFC 2045](#), November 1996.
- [RFC2046] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", [RFC 2046](#), November 1996.
- [RFC3501] Crispin, M., "INTERNET MESSAGE ACCESS PROTOCOL - VERSION 4rev1", [RFC 3501](#), March 2003.

Appendix A. Acknowledgements

Many thanks for input provided by Pete Resnick, James Galvin, Ned Freed, John Klensin and members of the EAI WG to help solidify this proposal.

Appendix B. Open Issues

Suggestion to change the utf-8-addr format from \-encoded Unicode to \-encoded UTF-8 as used in URIs.

Use a single syntax for I18N addresses in ORCPT/DSN instead of two (Chris)

Potential issue: an SMTP server can't deliver an EAI DSN to the next hop - need to use a 7bit encoding, downgrade or discard? Need to describe choices.

Tracker issue #1485: UTF8HDR 4.6/DSN: Choice of body part for transport of UTF8SMTP messages

Tracker issue #1483: SMTPEXT 2.7: Non-ASCII in response texts

Appendix C. Changes from -01

Cleaned up and tightened ABNF, in particular HEXPOINT.

Extended DSN report syntax to allow for localized version of diagnostic-code-field.

Added ABNF for the EAI DSN and EAI MDN.

Appendix D. Changes from -00

Added paragraph about use of 8bit Content-Transfer-Encoding for new message sub-types.

Updated the list of open issues.

Clarified that this document is targeted to become an Experimental RFC.

Made the EAI downgrade document a normative reference.

Updated ABNF for utf-8-address.

Authors' Addresses

Chris Newman
Sun Microsystems
3401 Centrelake Dr., Suite 410
Ontario, CA 91761
US

Email: chris.newman@sun.com

Alexey Melnikov (editor)
Isode Ltd
5 Castle Business Village
36 Station Road
Hampton, Middlesex TW12 2BX
UK

Email: Alexey.Melnikov@isode.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

