

UTF-8 Mail: Scenarios
draft-ietf-eai-scenarios-01

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 December 28, 2006.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes some scenarios in which one can imagine internationalized email addresses deployed, and tries to draw some conclusions about what's acceptable and what's not for users in those scenarios.

One possible set of extensions that can work in these scenarios is those described in the UTF8MAIL extension documents.

Requirements Language

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](#)].

Table of Contents

1.	Introduction	3
1.1.	Terminology	3
1.2.	User interface issues	3
1.3.	Ignored issues	4
2.	Important Scenarios	4
2.1.	Two i18mail users	4
2.2.	Three i18mail users	5
2.3.	i18mail mailing list	5
2.4.	One i18mail user sends to one ascii user	5
2.5.	An i18mail user sends to one ascii user and one i18mail user	6
2.6.	An i18mail user sends to a mailing list with a mix of users	6
2.7.	An i18mail user forwards to an ASCII user	7
3.	Other scenarios	7
3.1.	Two i18mail users, intermediate non-extended MTA	7
4.	IANA Considerations	7
5.	Security Considerations	8
6.	Acknowledgements	8
7.	Normative References	8
	Author's Address	9
	Intellectual Property and Copyright Statements	10

1. Introduction

With the advent of internationalized email addresses [ref], there is a very real risk that people using Internet email will have problems communicating that they did not have before. This document tries to sketch some of the scenarios, define what "proper" behaviour would be in the situations, and describe how this will constrain solutions to the "internationalized email" problem.

Because of the well known phenomenon that short documents get more review, the document tries to be as brief as possible, as long as that does not sacrifice clarity.

1.1. Terminology

Terminology is inherited from the UTF8SMTP framework [I-D.ietf-eai-framework] - in particular, the terms "ascii address", "internationalized email address" (EAI), "ascii user", "i18mail user", "message" and "mailing list" are used with the definitions from that document.

The term "UTF8SMTP" is used to refer to the particular solution proposed in that framework, while "i18n mail" refers to any solution that could conceivably satisfy these scenarios' requirements.

The pronoun "he" is used to indicate a human of indeterminate sex.

1.2. User interface issues

In internationalization, one of the thornier issues has always been user interfaces. In particular in this context, the ability to manipulate text strings (email addresses, in this case) in a script that the user does not have familiarity with is an issue.

A main purpose of i18mail is to allow users to avoid doing so when corresponding with users in their own language locale when that locale normally does not use ASCII - but unless we accept an Internet email community of many small fragments, the introduction of "local" characters into email addresses will cause users to be exposed to addresses that they have trouble recognizing, are unable to enter, and in fact may be unable to display; in some cases, even storing the addresses is an issue.

For instance, handling of right-to-left scripts like Arabic in environments used to left-to-right scripts like Thai can be a serious challenge.

In order to keep this document short, the following capabilities are

assumed:

- o An i18mail user is able to enter and display directly all characters of interest to him in his language locale.
- o An i18mail user is able to display all valid characters for EAI addresses, store them in an address book, and use "reply" without damaging the address.
- o If the i18mail solution requires keeping extra information around for an address in some cases, the i18mail user is capable of manipulating that information, including storing that information in his address book

One can imagine special circumstances where some of these do not represent an optimal solution (for instance, a Thai user may prefer to handle the ASCII address of an Arabic correspondent rather than his Arabic one), but this is an added complication, and is ignored for the moment.

1.3. Ignored issues

All the scenarios assume that all parties desire to communicate, that spam filters do not eat messages randomly, and that the mail service behaves according to specification. These are not tenable assumptions in the real world, but considering them would make this document much longer.

2. Important Scenarios

In the scenario descriptions below, A, B and C are i18mail users. X (and Y and Z if they need to occur) are ascii users. L is an i18n-aware mailing list; LA (not used yet) is a non-i18n-aware mailing list.

Apart from the messages being exchanged, and A knowing the addresses of the ones he sends mail to, which are presumed to be made known to A through some other method (business cards, web pages, mail from other users and directories are some examples), there is no communication required between the users.

2.1. Two i18mail users

One i18mail user (A) sends a message to another i18mail user (B), and desires to use his EAI address. The recipient replies to the message.

Requirement: The message must arrive at the recipient. The reply must arrive at the sender. The email addresses visible to the sender and recipient must be the EAI addresses.

2.2. Three i18mail users

As above, but A sends his message to both B and C. Both reply to all the recipients listed in the message.

Requirement: As above - B and C must get the message, A and C must get the reply from B, A and B must get the reply from C. The email addresses visible to A, B and C must all be the EAI addresses.

2.3. i18mail mailing list

A sends his message to L, a mailing list, which has subscribers B and C. Both reply to the mailing list. The mailing list is i18n aware.

Requirement: As above - all messages arrive, with EAI addresses preserved for all 3 users.

2.4. One i18mail user sends to one ascii user

A, an i18mail user, sends to X, an ascii user. X replies.

In this scenario, it is a given that A, the sender, has to have some facility for handling ASCII; he has to at least be able to display and enter an ASCII address.

Precondition: A has to have an ascii address.

Requirement: There must be an algorithmic series of steps that A can follow in order to get a message to X, and where X's reply gets back to A.

There is no requirement that X sees the EAI address of A, or that the address of A that X sees be one that A knows about beforehand; the requirement is that the messages get there. This non-requirement applies to all the following cases too.

Examples of ways this could happen:

- o Magic happens in the network: A's message gets converted in the network to a format acceptable to X. This may require A to include extra information with the message to help the conversion process - and may be impossible to do for the general case.

- o Sender selection: A's i18mail message gets bounced in the network, and the reception of the error report causes A to resend the message using a format acceptable to X.
- o Conversion at destination) A's message gets accepted, and X has facilities available to convert it into an acceptable form for X, including deriving a valid ASCII address for A. This would require knowledge of i18n at X's site, but not necessarily in X's user agent.

This is NOT an exhaustive list, and is NOT part of the requirements of the scenarios. A given protocol for i18mail will in turn impose new requirements on the scenarios - for instance, if extra information is included with the message, a user interface may need to exist to allow the sender to manipulate this information.

2.5. An i18mail user sends to one ascii user and one i18mail user

In this scenario, A sends to B and X; both reply.

Precondition: A and B have to have valid ASCII addresses.

Requirement: Through some series of steps, A must be able to get a message to both B and X; through some series of steps, B and X must be able to reply to each other and to A. X must not require information outside of what is included in the message to get a message to B.

Possible non-requirements (for discussion):

- o Maybe the messages to B and X don't need to be exactly the same.
- o Maybe B doesn't need to see or use A's EAI address when he's replying to A and X.
- o Maybe X doesn't need to see A's address exactly the same on the message from A and the reply from B.

2.6. An i18mail user sends to a mailing list with a mix of users

In this scenario, A sends to L, and L has B and X as subscribers. B and X reply.

Requirement: Messages get there. A will not have to know anything about X in order to make the messages go through.

Notes and non-requirements:

- o It may be acceptable for A to have to treat L as if L was an ASCII mailing list (LA)
- o It may be acceptable for B to see A's ASCII address, not his EAI address
- o How one can transition between this and the scenario of [Section 2.3](#) is unclear.

[2.7.](#) An i18mail user forwards to an ASCII user

In this scenario, A sends to B, and B forwards the message as a MIME attachment to X.

Precondition: B has valid ASCII addresses.

Requirement: The message from B to X should arrive whether or not A's address is usable by X.

Desirable property: If A's address is downgradable, it should be usable by X for generating a message.

[3.](#) Other scenarios

This section collects scenarios that have been discussed, but where there is no WG consensus on whether or not they are important enough to influence the design of UTF8SMTP.

[3.1.](#) Two i18mail users, intermediate non-extended MTA

In this scenario, A sends mail to B through an MTA that does not support i18mail extensions.

Requirement: Mail arrives.

Desirable property: B can see A's I18mail address in his user interface, and use that to reply.

The reason this may not be an important scenario is that due to the largely end-to-end nature of SMTP, if the end users have upgraded their systems, there should be very little reason to go via a non-upgraded MTA.

[4.](#) IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

5. Security Considerations

Security issues are deliberately overlooked in order to reduce the size of the document.

6. Acknowledgements

7. Normative References

[I-D.ietf-eai-framework]

Klensin, J. and Y. Ko, "Overview and Framework for Internationalized Email", [draft-ietf-eai-framework-00](#) (work in progress), May 2006.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

Author's Address

Harald Tveit Alvestrand
Google

Intellectual Property Statement

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.

Disclaimer of Validity

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 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.

Copyright Statement

Copyright (C) The Internet Society (2006). 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.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

