LAMPS
Internet-Draft
Intended status: Standards Track
Expires: September 13, 2017

A. Melnikov, Ed.
Isode Ltd
W. Chuang, Ed.
Google, Inc.

March 12, 2017

Internationalized Email Addresses in X.509 certificates
draft-ietf-lamps-eai-addresses-08

Abstract

This document defines a new name form for inclusion in the otherName field of an X.509 Subject Alternative Name and Issuer Alternate Name extension that allows a certificate subject to be associated with an Internationalized Email Address.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on September 13, 2017.

Copyright Notice

Table of Contents

1. Introduction

[RFC5280] defines rfc822Name subjectAltName choice for representing [RFC5321] email addresses. This form is restricted to a subset of US-ASCII characters and thus can't be used to represent Internationalized Email addresses [RFC6531]. To facilitate use of these Internationalized Email addresses with X.509 certificates, this document specifies a new name form in otherName so that subjectAltName and issuerAltName can carry them. In addition this document calls for all email address domain in X.509 certificates to conform to IDNA2008 [RFC5890].

2. Conventions Used in This Document

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 [RFC2119].

The formal syntax use the Augmented Backus-Naur Form (ABNF) [RFC5234] notation.

3. Name Definitions

The GeneralName structure is defined in [RFC5280], and supports many different names forms including otherName for extensibility. This section specifies the SmtpUTF8Name name form of otherName, so that Internationalized Email addresses can appear in the subjectAltName of

a certificate, the issuerAltName of a certificate, or anywhere else that GeneralName is used.

```
id-on-SmtpUTF8Name OBJECT IDENTIFIER ::= { id-on 9 }

SmtpUTF8Name ::= UTF8String (SIZE (1..MAX))
```

When the subjectAltName (or issuerAltName) extension contains an Internationalized Email address, the address MUST be stored in the SmtpUTF8Name name form of otherName. The format of SmtpUTF8Name is defined as the ABNF rule SmtpUTF8Mailbox. SmtpUTF8Mailbox is a modified version of the Internationalized Mailbox which was defined in Section 3.3 of [RFC6531] which was itself derived from SMTP Mailbox from Section 4.1.2 of [RFC5321]. [RFC6531] defines the following ABNF rules for Mailbox whose parts are modified for internationalization: <Local-part>, <Dot-string>, <Quoted-string>, <QcontentSMTP>, <Domain>, and <Atom>. In particular, <Local-part> was updated to also support UTF8-non-ascii. UTF8-non-ascii was described by Section 3.1 of [RFC6532]. Also, sub-domain was extended to support U-label, as defined in [RFC5890].

This document further refines Internationalized [RFC6531] Mailbox ABNF rules and calls this SmtpUTF8Mailbox. In SmtpUTF8Mailbox, sub-domain that encode non-ASCII characters SHALL use U-label Unicode native character labels and MUST NOT use A-label [RFC5890]. This restriction prevents having to determine which label encoding A- or U-label is present in the Domain. As per Section 2.3.2.1 of [RFC5890], U-label use UTF-8 [RFC3629] with Normalization Form C and other properties specified there. In SmtpUTF8Mailbox, sub-domain that encode ASCII character labels SHALL use NR-LDH restrictions as specified by section 2.3.1 of [RFC5890] and SHALL be restricted to lower case letters. One suggested approach to apply these sub-domains restriction is to restrict sub-domain so that labels not start with two letters followed by two hyphen-minus characters. Consistent with the treatment of rfc822Name in [RFC5280], SmtpUTF8Name is an envelope <Mailbox> and has no phrase (such as a

common name) before it, has no comment (text surrounded in parentheses) after it, and is not surrounded by "<" and ">".

In the context of building name constraint as needed by [RFC5280], the SmtpUTF8Mailbox rules are modified to allow partial productions to allow for additional forms required by Section 6. Name constraints may specify a complete email address, host name, or domain. This means that the local-part may be missing, and domain partially specified.

SmtpUTF8Name is encoded as UTF8String. The UTF8String encoding MUST NOT contain a Byte-Order- Mark (BOM) [RFC3629] to aid consistency across implementations particularly for comparison.

4. IDNA2008

To facilitate comparison between email addresses, all email address domain in X.509 certificates MUST conform to IDNA2008 [RFC5890]. Otherwise non-conforming email address domains introduces the possibility of conversion errors between alternate forms. This applies to SmtpUTF8Mailbox and rfc822Name in subjectAltName, issuerAltName and anywhere else that GeneralName is used.

5. Matching of Internationalized Email Addresses in X.509 certificates

In equivalence comparison with SmtpUTF8Name, there may be some setup work to enable the comparison i.e. processing of the SmtpUTF8Name content or the email address that is being compared against. The process for setup for comparing with SmtpUTF8Name is split into domain steps and local- part steps. The comparison form for local-part always is UTF-8. The comparison form for domain depends on context. While some contexts such as certificate path validation in [RFC5280] specify transforming domain to A-label, this document RECOMMENDS transforming to UTF-8 U-label instead. This reduces the likelihood of errors by reducing conversions as more implementations natively support U-label domains.

Comparison of two SmtpUTF8Name is straightforward with no setup work needed. They are considered equivalent if there is an exact octet-

for-octet match. Comparison with other email address forms such as Internationalized email address or rfc822Name requires additional setup steps. Domain setup is particularly important for forms that may contain A- or U-label such as International email address, or A-label only forms such as rfc822Name. This document specifies the process to transform the domain to U-label. (To convert the domain to A-label, follow the process specified in section 7.5 and 7.2 in [RFC5280]) The first step is to detect A-label by using section 5.1 of [RFC5891]. Next if necessary, transform the A-label to U-label Unicode as specified in section 5.2 of [RFC5891]. Finally if necessary convert the Unicode to UTF-8 as specified in section 3 of [RFC3629]. For ASCII NR-LDH labels, upper case letters are converted to lower case letters. In setup for SmtpUTF8Mailbox, the email address local-part MUST conform to the requirements of [RFC6530] and [RFC6531], including being a string in UTF-8 form. In particular, the local-part MUST NOT be transformed in any way, such as by doing case folding or normalization of any kind. The <Local-part> part of an Internationalized email address is already in UTF-8. For rfc822Name the local-part, which is IA5String (ASCII), trivially maps

to UTF-8 without change. Once setup is complete, they are again compared octet-for-octet.

To summarize non-normatively, the comparison steps including setup are:

1. If the domain contains A-labels, transform them to U-label.

2. If the domain contains ASCII NR-LDH labels, lowercase them.

3. Ensure local-part is UTF-8.

4. Compare strings octet-for-octet for equivalence.

This specification expressly does not define any wildcards characters and SmtpUTF8Name comparison implementations MUST NOT interpret any character as wildcards. Instead, to specify multiple email addresses through SmtpUTF8Name, the certificate SHOULD use multiple subjectAltNames or issuerAltNames to explicitly carry those email addresses.

6. Name constraints in path validation

This section defines use of SmtpUTF8Name name for name constraints. The format for SmtpUTF8Name in name constraints is identical to the use in subjectAltName as specified in Section 3 with the extension as noted there for partial productions.

Constraint comparison on complete email address with SmtpUTF8Name name uses the matching procedure defined by Section 5. As with rfc822Name name constraints as specified in Section 4.2.1.10 of [RFC5280], SmtpUTF8Name name can specify a particular mailbox, all addresses at a host, or all mailboxes in a domain by specifying the complete email address, a host name, or a domain. Name constraint comparisons in the context of [RFC5280] that are specified with SmtpUTF8Name name are only done on the subjectAltName SmtpUTF8Name name and not on other forms. Similarly rfc822Name name constraints do not apply to subjectAltName SmtpUTF8Name name. This imposes requirements on the certificate issuer as described next.

When name constraints are used with SmtpUTF8Name subject alternative names, the constraints are specified by the following changes to the path validator to prevent bypass of the name constraints. The email address path validator in Section 6 of [RFC5280] is modified to consider:

1. When neither rfc822Name nor SmtpUTF8Name name constraints are present in any issuer CA certificate, then path validation does

not add restrictions on children certificates with rfc822Name or SmtpUTF8Name subject alternative names. That is any combination of rfc822Name or SmtpUTF8Name subject alternative names may be present.

2. If issuer CA certificates contain only rfc822Name name constraints, then those constraints apply to rfc822Name subject alternative name in children certificates. SmtpUTF8Name subject alternative name are prohibited in those same certificates, that is those certificates MUST be rejected by the path verifier.

3. When both rfc822Name and SmtpUTF8Name name constraints are present in all issuer CA certificates that have either form, then the path verifier applies the constraint of the subject alternative name form in children certificates. This allows any

combination of rfc822Name or SmtpUTF8Name subject alternative names to be present and implies that the issuer has applied appropriate name constraints. While commonly the alternative forms will be equivalent, they need not be, as the forms can represent features not present in its counterpart. One instance of this is when the issuer wants to name constrain domain or hostname using the rules of a particular form.

4. If some issuer CA certificates contain only SmtpUTF8Name name constraints, then those are at risk of bypass with rfc822Name subject alternative names when processed by legacy verifiers. To prevent this, issuers MUST also publish rfc822Name name constraint that prevent those bypasses. This occurs when both rfc822Name and SmtpUTF8Name constraint forms can represents the same host, domain or email address, and both are needed. Even when the constraints are asymmetric such as when the issuer wishes to constrain an email address with an UTF-8 local part, a non empty rfc822Name name constraint may be needed if there isn't one already so that the path verifier initializes correctly.

When both name constraints are present, the contents depends on the usage. If the issuer desires to represent the same NR-LDH host or domain, then it is the same string in both rfc822Name and SmtpUTF8Name. If the host or domain labels contain UTF-8, then the labels may be used directly in SmtpUTF8Name noting the restriction in Section 5 and transformed to A-label for rfc822Name using the process described in [RFC5280]. Email addresses that use ASCII local-part use the same processing procedures for host or domain.

If the issuer wishes to represent the name constraint asymmetrically, with either rfc822Name or SmtpUTF8Name to respectively represent some A-label or U-label in the domain or host, the alternate name constraint form must still be present. If nothing needs be

represented by the alternate form, then empty name constraint can described by the "invalid" TLD that helps initialize the name constraint path validation set. Or alternatively it may be omitted if some other name constraint pair, provides a name constraint of that form. In particular this initialization may be needed when SmtpUTF8Name is used to represent an email address name constraint with an UTF-8 local-part and rfc822Name cannot represent such a email address constraint.

The name constraint requirement with SmtpUTF8Name subject alternative name is illustrated in the non-normative diagram Figure 1 with several examples. (3a) shows an issuer constraining a NR-LDH hostname with rfc822Name and SmtpUTF8Name so that they can issue ASCII and UTF-8 local-name email addresses certificates. (3b) shows an issuer constraining a hostname containing a non-ASCII label for u+5C0Fu+5B66 (elementary school). (3c) demonstrates that a hostname constraint with an rfc822Name is distinguishable from its SmtpUTF8Name constraint, and that only the rfc822Name form is permitted. No 'invalid' SmtpUTF8Name constraint is needed since other SmtpUTF8Name constraints are present. (3d) similarly demonstrates this capability to restrict a name constraint to SmtpUTF8Name only. (3e) shows that a non-ASCII local- part email address can also be constrained to be permitted using SmtpUTF8Name. It too does not need an 'invalid' rfc822Name as other rfc822Name constrains are present. Diagram Figure 2 illustrates (non-normatively) a different certificate chain that does need the 'invalid' name constraint. (3f) constrains a non-ASCII local-part email address using a SmtpUTF8Name name constraint but requires a rfc822Name 'invalid' constraint because it lacks any other rfc822Name constraints needed to initialize the name constraint path verification. The next non-normative diagram Figure 3 illustrates legacy name constraints that contrasts the changes this document specifies. The legacy approach (2) has only a single rfc822Name name email address name constraint.

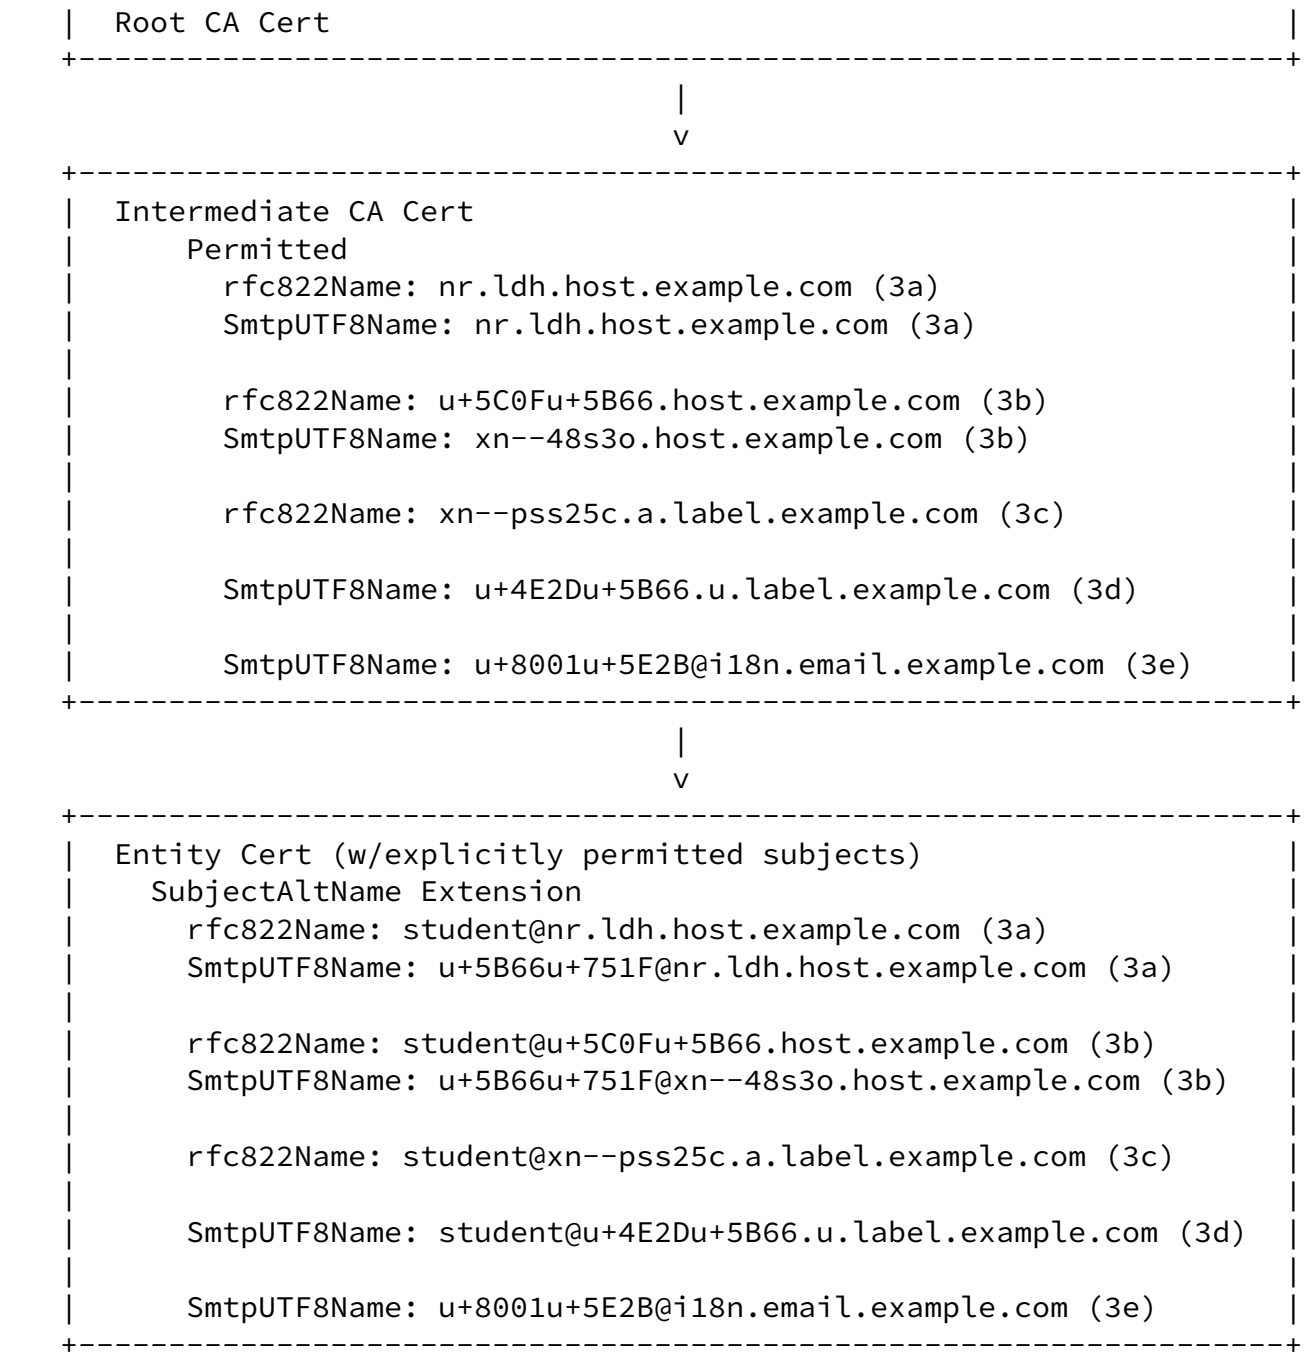

Name constraints with SmtpUTF8Name and rfc822Name

Figure 1

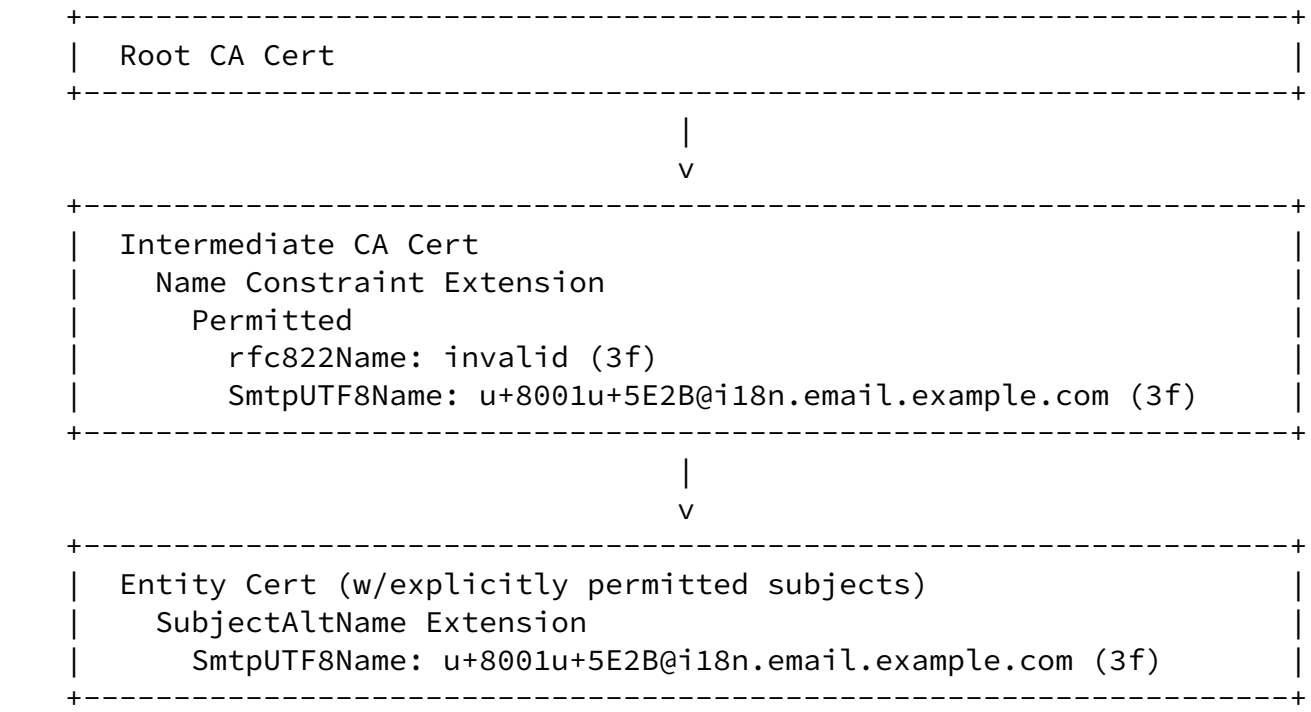

Name constraints with SmtpUTF8Name email address and empty rfc822Name

Figure 2

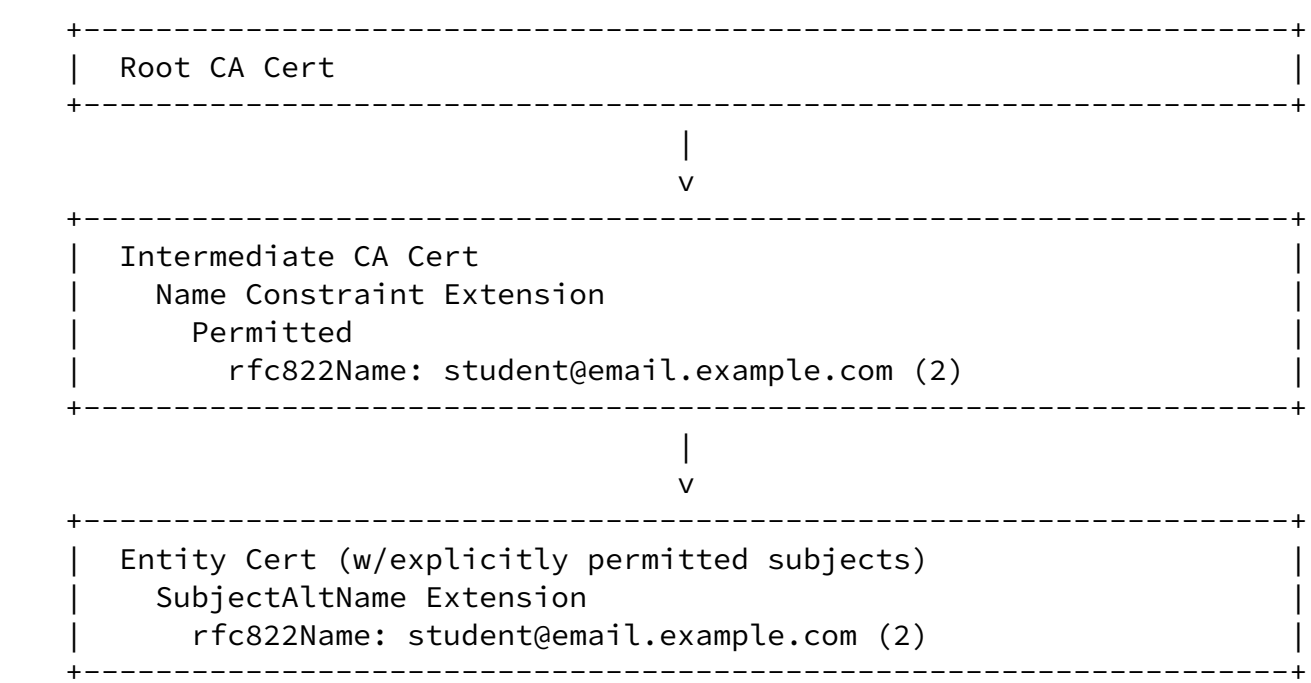

Legacy name constraints with rfc822Name

Figure 3

7. Deployment Considerations

For email addresses whose local-part is ASCII it may be more reasonable to continue using rfc822Name instead of SmtpUTF8Name. The use of rfc822Name rather than SmtpUTF8Name is currently more likely to be supported. Also use of SmtpUTF8Name incurs higher byte representation overhead due to encoding with otherName and the additional OID needed. This may be offset if domain requires non-ASCII characters as SmtpUTF8Name supports U-label whereas rfc822Name supports A-label. This document RECOMMENDS using SmtpUTF8Name when local-part contains non-ASCII characters, and otherwise rfc822Name.

8. Security Considerations

Use for SmtpUTF8Name for certificate subjectAltName (and issuerAltName) will incur many of the same security considerations of Section 8 in [RFC5280] but is further complicated by permitting non-ASCII characters in the email address local-part. This complication, as mentioned in Section 4.4 of [RFC5890] and in Section 4 of [RFC6532], is that use of Unicode introduces the risk of visually similar and identical characters which can be exploited to deceive the recipient. The former document references some means to mitigate against these attacks.

9. IANA Considerations

in Section Section 3 and the ASN.1 module identifier defined in Section Appendix A. IANA is kindly requested to make the following assignments for:

The LAMPS-EaiAddresses-2016 ASN.1 module in the "SMI Security for PKIX Module Identifier" registry (1.3.6.1.5.5.7.0).

The SmtpUTF8Name otherName in the "PKIX Other Name Forms" registry (1.3.6.1.5.5.7.8).

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <http://www.rfc-editor.org/info/rfc3629>.

[RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <http://www.rfc-editor.org/info/rfc5234>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <http://www.rfc-editor.org/info/rfc5280>.

[RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <http://www.rfc-editor.org/info/rfc5321>.

[RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <http://www.rfc-editor.org/info/rfc5890>.

[RFC5891] Klensin, J., "Internationalized Domain Names in Applications (IDNA): Protocol", RFC 5891, DOI 10.17487/RFC5891, August 2010, <http://www.rfc-editor.org/info/rfc5891>.

[RFC6530] Klensin, J. and Y. Ko, "Overview and Framework for Internationalized Email", RFC 6530, DOI 10.17487/RFC6530, February 2012, <http://www.rfc-editor.org/info/rfc6530>.

[RFC6531] Yao, J. and W. Mao, "SMTP Extension for Internationalized Email", RFC 6531, DOI 10.17487/RFC6531, February 2012, <http://www.rfc-editor.org/info/rfc6531>.

[RFC6532] Yang, A., Steele, S., and N. Freed, "Internationalized Email Headers", RFC 6532, DOI 10.17487/RFC6532, February 2012, <http://www.rfc-editor.org/info/rfc6532>.

10.2. Informative References

[RFC5912] Hoffman, P. and J. Schaad, "New ASN.1 Modules for the Public Key Infrastructure Using X.509 (PKIX)", RFC 5912, DOI 10.17487/RFC5912, June 2010, <http://www.rfc-editor.org/info/rfc5912>.

Appendix A. ASN.1 Module

The following ASN.1 module normatively specifies the SmtpUTF8Name structure. This specification uses the ASN.1 definitions from [RFC5912] with the 2002 ASN.1 notation used in that document. [RFC5912] updates normative documents using older ASN.1 notation.

```
LAMPS-EaiAddresses-2016
  { iso(1) identified-organization(3) dod(6)
    internet(1) security(5) mechanisms(5) pkix(7) id-mod(0)
    id-mod-lamps-eai-addresses-2016(TBD) }

DEFINITIONS IMPLICIT TAGS ::=
BEGIN

IMPORTS
  OTHER-NAME
  FROM PKIX1Implicit-2009
    { iso(1) identified-organization(3) dod(6) internet(1) security(5)
    mechanisms(5) pkix(7) id-mod(0) id-mod-pkix1-implicit-02(59) }

  id-pkix
  FROM PKIX1Explicit-2009
    { iso(1) identified-organization(3) dod(6) internet(1) security(5)
    mechanisms(5) pkix(7) id-mod(0) id-mod-pkix1-explicit-02(51) } ;
```

```
--
-- otherName carries additional name types for subjectAltName,
-- issuerAltName, and other uses of GeneralNames.
--

  id-on OBJECT IDENTIFIER ::= { id-pkix 8 }

  SmtpUtf8OtherNames OTHER-NAME ::= { on-SmtpUTF8Name, ... }

  on-SmtpUTF8Name OTHER-NAME ::= {
      SmtpUTF8Name IDENTIFIED BY id-on-SmtpUTF8Name
  }

  id-on-SmtpUTF8Name OBJECT IDENTIFIER ::= { id-on 9 }

  SmtpUTF8Name ::= UTF8String (SIZE (1..MAX))

END
```

Figure 4

Appendix B. Example of SmtpUTF8Name

This non-normative example demonstrates using SmtpUTF8Name as an otherName in GeneralName to encode the email address "u+8001u+5E2B@example.com".

The hexadecimal DER encoding of the email address is:

```
A022060A 2B060105 05070012 0809A014 0C12E880 81E5B8AB 40657861
6D706C65 2E636F6D
```

The text decoding is:

```
  0  34: [0] {
  2  10:   OBJECT IDENTIFIER '1 3 6 1 5 5 7 0 18 8 9'
 14  20:   [0] {
 16  18:     UTF8String '..@example.com'
       :     }
       :   }
```

Figure 5

The example was encoded on the OSS Nokalva ASN.1 Playground and the above text decoding is an output of Peter Gutmann's "dumpasn1" program.

Appendix C. Acknowledgements

Thank you to Magnus Nystrom for motivating this document. Thanks to Russ Housley, Nicolas Lidzborski, Laetitia Baudoin, Ryan Sleevi, Sean Leonard, Sean Turner, John Levine, and Patrik Falstrom for their feedback. Also special thanks to John Klensin for his valuable input on internationalization, Unicode and ABNF formatting, to Jim Schaad for his help with the ASN.1 example and his helpful feedback, and to Viktor Dukhovni for his help with name constraints.

Authors' Addresses

Alexey Melnikov (editor)
Isode Ltd
14 Castle Mews
Hampton, Middlesex TW12 2NP
UK

Email: Alexey.Melnikov@isode.com

Weihaw Chuang (editor)
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: weihaw@google.com