LAMPS
Internet-Draft
Intended status: Standards Track
Expires: May 3, 2017

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

October 30, 2016

Internationalized Email Addresses in X.509 certificates
draft-ietf-lamps-eai-addresses-01

Abstract

This document defines a new name form for inclusion in the otherName field of an X.509 Subject Alternative 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 May 3, 2017.

Table of Contents

1. Introduction

[RFC5280] defines rfc822Name subjectAltName choice for representing [RFC5322] 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 fascilitate 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.

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 is defined in Section 3.3 of [RFC6531] which is 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 is described by Section 3.1 of [RFC6532]. Also sub-domain is 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 solely ASCII character labels SHALL use NR-LDH restrictions as specified by section 2.3.1 of [RFC5890]. Note that a smtputf8Mailbox 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 5. 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. 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 to A-label, this document RECOMMENDS transforming to UTF-8 U-label even in place of those other specifications. As more implementations natively support U-label domain, requiring U-label reduces conversions required, which then reduces likelihood of errors caused by bugs in implementation.

Comparison of two smtputf8Name can be straightforward. No setup work is needed and it can be an octet for octet comparison. For other email address forms such as Internationalized email address or rfc822Name, the comparison requires additional setup to convert the format for comparison. 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 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]. In setup for smtputf8Mailbox, the email address local-part MUST be converted to UTF-8 if it is not already. The <Local-part> part of an Internationalized email address is already in UTF-8. For the rfc822Name local-part is IA5String (ASCII), and conversion to UTF-8 is trivial since ASCII octets maps to UTF-8 without change. Once the setup is completed, comparison is an octet for octet comparison.

This specification expressly does not define any wildcards characters and smtputf8Name comparison implementations MUST NOT interpret any character as wildcards. Instead, to specify multiple specifying multiple email addresses through smtputf8Name, the certificate should use multiple subjectAltNames or issuerAltNames to explicitly carry those email addresses.

5. 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 4. 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 [RFC5280] is specified with smtputf8Name name are only done on the subjectAltName (and issuerAltName) smtputf8Name name, and says nothing more about constaints on other email address forms such as rfc822Name. Consequently it may be necessary to include other name constraints such as rfc822Name in addition to smtputf8Name to constrain all potential email addresses. For example a domain with both ascii and non-ascii local-part email addresses may require both rfc822Name and smtputf8Name name constraints. This can be illustrated in the following non-normative diagram Figure 1 which shows a name constraint set in the intermediate CA certificate, which then applies to the children entity certificates. Note that a constraint on rfc822Name does not apply to smtputf8Name and vice versa.

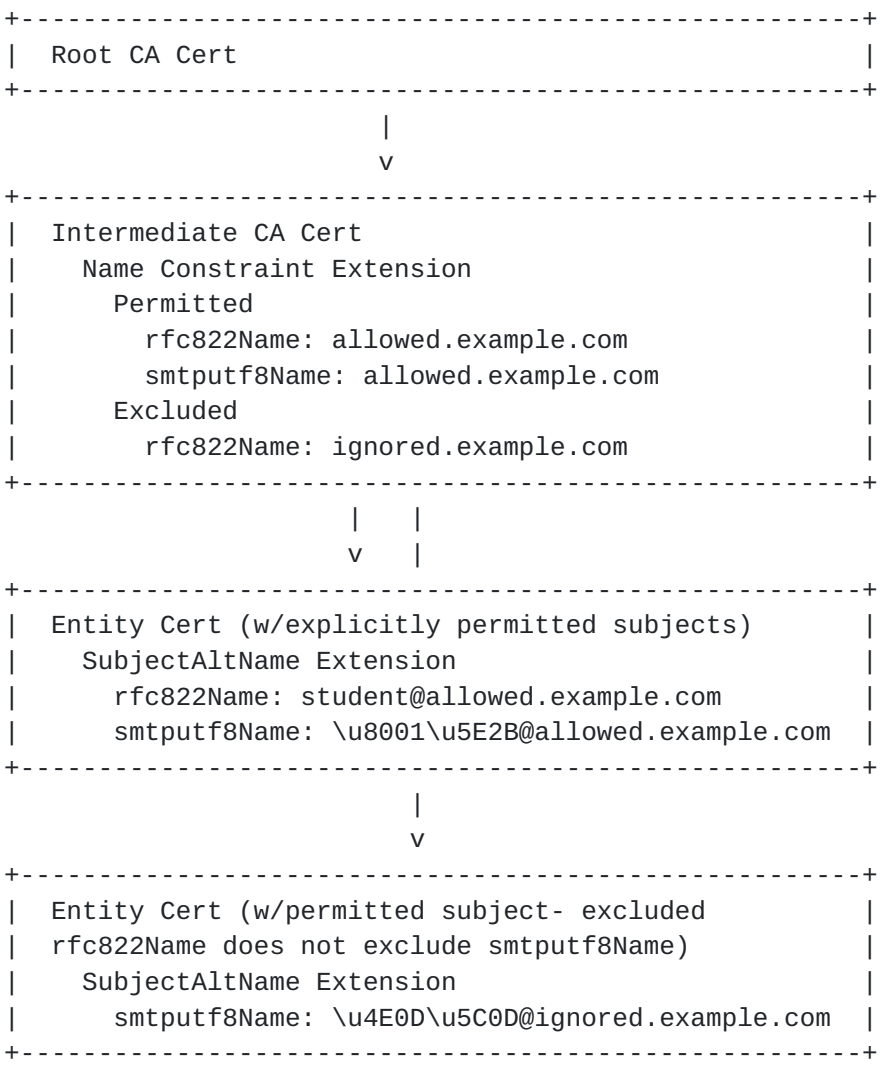

Figure 1

6. Resource Considerations

For email addresses whose local-part is ASCII it may be more reasonable to continue using rfc822Name instead of smtputf8Name. Use of smtputf8Name incurs higher byte representation overhead due to encoding with otherName and the additional OID needed. This document RECOMMENDS using smtputf8Name when local-part contains non-ASCII characters, and otherwise rfc822Name.

7. Security Considerations

Use for smtputf8Name for certificate subjectAltName (and issuerAltName) will incur many of the same security considerations of Section 8 in [RFC5280] but further complicated by permitting non-ASCII characters in the email address local-part. As mentioned in

Section 4.4 of [RFC5890] and in Section 4 of [RFC6532] Unicode introduces the risk for visually similar characters which can be exploited to deceive the recipient. The former document references some means to mitigate against these attacks.

8. IANA Considerations

This document makes use of object identifiers for the smtputf8Name defined 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).

9. References

9.1. 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>.

[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>.

[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>.

9.2. Informative References

[RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, DOI 10.17487/RFC5322, October 2008, <http://www.rfc-editor.org/info/rfc5322>.

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.

```
  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(88) }

  DEFINITIONS IMPLICIT TAGS ::=
  BEGIN

  IMPORTS

    id-pkix  OBJECT IDENTIFIER  ::=
        {iso(1) identified-organization(3) dod(6) internet(1) security(5)
        mechanisms(5) pkix(7)}

  --
  -- otherName carries additional name types for subjectAltName, issuerAltName,
  -- and other uses of GeneralNames.
  --
  -- Note that the LAMPS-EaiAddresses-2016 module and id-on-smtputf8Name OID
  -- uses example IANA numbers i.e. are non-normative.
  --

    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 2

Appendix B. Acknowledgements

Thank you to Magnus Nystrom for motivating this document. Thanks to Russ Housley, Nicolas Lidzborski, Laetitia Baudoin, Ryan Sleevi, Sean Leonard, Sean Turner, and Jim Schaad for their feedback. Also thanks to John Klensin for his valuable input on internationalization, Unicode and ABNF formatting.

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