Internet Engineering Task Force (IETF)
Request for Comments: 9549
Obsoletes: 8399
Updates: 5280
Category: Standards Track
ISSN: 2070-1721

R. Housley
Vigil Security

March 2024

Internationalization Updates to RFC 5280

Abstract

The updates to RFC 5280 described in this document provide alignment with the 2008 specification for Internationalized Domain Names (IDNs) and includes support for internationalized email addresses in X.509 certificates. The updates ensure that name constraints for email addresses that contain only ASCII characters and internationalized email addresses are handled in the same manner. This document obsoletes RFC 8399.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9549.

Table of Contents

1. Introduction

This document updates the Introduction in Section 1, the Name Constraints certificate extension discussion in Section 4.2.1.10, and the Processing Rules for Internationalized Names in Section 7 of RFC 5280 [RFC5280] to provide alignment with the 2008 specification for Internationalized Domain Names (IDNs) and includes support for internationalized email addresses in X.509 certificates.

An IDN in Unicode (native character) form contains at least one U-label [RFC5890]. IDNs are carried in certificates in ACE-encoded form. That is, all U-labels within an IDN are converted to A-labels. Conversion of a U-label to an A-label is described in [RFC5891].

The GeneralName structure supports many different name forms, including otherName for extensibility. RFC 8398 [RFC8398] specifies the SmtpUTF8Mailbox for internationalized email addresses.

Note that Internationalized Domain Names in Applications specifications published in 2003 (IDNA2003) [RFC3490] and 2008 (IDNA2008) [RFC5890] both refer to the Punycode algorithm for conversion [RFC3492].

Note that characters in the Unicode Category "Symbol, Other" (So) are specifically not included in IDNA2003 [RFC3490] or IDNA2008 [RFC5890]; the derived property values for characters in this category are calculated as DISALLOWED. Thus, some characters that are allowed under the Unicode IDNA Compatibility Processing [UTS46] are not allowed under this specification. For instance, .example, which contains the Unicode character U+1F0A1 (BLACK CHESS KING), results in a failure under this specification, but it becomes xn--45h.example under [UTS46].

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Changes since RFC 8399

In some cases, [RFC8399] required conversion of A-labels to U-labels in order to process name constraints for internationalized email addresses. This led to implementation complexity and at least two security vulnerabilities. One summary of the vulnerabilities can be

found in [DDHQ]. Now, all IDNs are carried and processed as A-labels.

The Introduction provides a warning to implementers about the handling of characters in the Unicode Category "Symbol, Other" (So), which includes emoji characters.

2. Updates to RFC 5280

This section provides updates to several paragraphs of [RFC5280]. For clarity, if the entire section is not replaced, then the original text and the replacement text are shown.

2.1. Update in the Introduction (Section 1)

This update provides references for IDNA2008.

OLD

| * Enhanced support for internationalized names is specified in
| Section 7, with rules for encoding and comparing
| Internationalized Domain Names, Internationalized Resource
| Identifiers (IRIs), and distinguished names. These rules are
| aligned with comparison rules established in current RFCs,
| including [RFC3490], [RFC3987], and [RFC4518].

NEW

| * Enhanced support for internationalized names is specified in
| Section 7, with rules for encoding and comparing
| Internationalized Domain Names, Internationalized Resource
| Identifiers (IRIs), and distinguished names. These rules are
| aligned with comparison rules established in current RFCs,
| including [RFC3987], [RFC4518], [RFC5890], and [RFC5891].

2.2. Update in Name Constraints (Section 4.2.1.10)

This update removes the ability to include constraints for a particular mailbox. This capability was not used, and removing it allows name constraints to apply to email addresses in rfc822Name and SmtpUTF8Mailbox [RFC8398] within otherName.

OLD

| A name constraint for Internet mail addresses MAY specify a
| particular mailbox, all addresses at a particular host, or all
| mailboxes in a domain. To indicate a particular mailbox, the
| constraint is the complete mail address. For example,
| "root@example.com" indicates the root mailbox on the host
| "example.com". To indicate all Internet mail addresses on a
| particular host, the constraint is specified as the host name.
| For example, the constraint "example.com" is satisfied by any mail
| address at the host "example.com". To specify any address within
| a domain, the constraint is specified with a leading period (as
| with URIs). For example, ".example.com" indicates all the
| Internet mail addresses in the domain "example.com", but not
| Internet mail addresses on the host "example.com".

NEW

| A name constraint for Internet mail addresses MAY specify all
| addresses at a particular host or all mailboxes in a domain. To
| indicate all Internet mail addresses on a particular host, the
| constraint is specified as the host name. For example, the
| constraint "example.com" is satisfied by any mail address at the
| host "example.com". To specify any address within a domain, the

| constraint is specified with a leading period (as with URIs). For
| example, ".example.com" indicates all the Internet mail addresses
| in the domain "example.com" but not Internet mail addresses on the
| host "example.com".

2.3. Update in IDNs in GeneralName (Section 7.2)

This update aligns with IDNA2008. Since all of Section 7.2 of [RFC5280] is replaced, the OLD text is not provided.

NEW

| Internationalized Domain Names (IDNs) may be included in
| certificates and CRLs in the subjectAltName and issuerAltName
| extensions, name constraints extension, authority information
| access extension, subject information access extension, CRL
| distribution points extension, and issuing distribution point
| extension. Each of these extensions uses the GeneralName type;
| one choice in GeneralName is the dNSName field, which is defined
| as type IA5String.
|
| IA5String is limited to the set of ASCII characters. To
| accommodate IDNs, U-labels are converted to A-labels. The A-label
| is the encoding of the U-label according to the Punycode algorithm
| [RFC3492] with the ACE prefix "xn--" added at the beginning of the
| string.
|
| When comparing DNS names for equality, conforming implementations
| MUST perform a case-insensitive exact match on the entire DNS
| name. When evaluating name constraints, conforming
| implementations MUST perform a case-insensitive exact match on a
| label-by-label basis. As noted in Section 4.2.1.10, any DNS name
| that may be constructed by adding labels to the left-hand side of
| the domain name given as the constraint is considered to fall
| within the indicated subtree.
|
| Implementations that have a user interface SHOULD convert IDNs to
| Unicode for display. Specifically, conforming implementations
| convert A-labels to U-labels for display purposes.
|
| Implementation consideration: There are increased memory
| requirements for IDNs. An IDN ACE label will begin with the four
| additional characters "xn--", and an IDN can require as many as
| five ASCII characters to specify a single international character.

2.4. Update in IDNs in Distinguished Names (Section 7.3)

This update aligns with IDNA2008.

OLD

| Domain Names may also be represented as distinguished names using
| domain components in the subject field, the issuer field, the
| subjectAltName extension, or the issuerAltName extension. As with
| the dNSName in the GeneralName type, the value of this attribute
| is defined as an IA5String. Each domainComponent attribute
| represents a single label. To represent a label from an IDN in
| the distinguished name, the implementation MUST perform the
| "ToASCII" label conversion specified in Section 4.1 of RFC 3490.
| The label SHALL be considered a "stored string". That is, the
| AllowUnassigned flag SHALL NOT be set.

NEW

| Domain names may also be represented as distinguished names using
| domain components in the subject field, the issuer field, the

| subjectAltName extension, or the issuerAltName extension. As with
| the dNSName in the GeneralName type, the value of this attribute
| is defined as an IA5String. Each domainComponent attribute
| represents a single label. To represent a label from an IDN in
| the distinguished name, the implementation MUST convert all
| U-labels to A-labels.

2.5. Update in Internationalized Electronic Mail Addresses (Section 7.5)

This update aligns with IDNA2008 and [RFC8398]. Since all of Section 7.5 of [RFC5280] is replaced, the OLD text is not provided.

NEW

| Electronic Mail addresses may be included in certificates and CRLs
| in the subjectAltName and issuerAltName extensions, name
| constraints extension, authority information access extension,
| subject information access extension, issuing distribution point
| extension, or CRL distribution points extension. Each of these
| extensions uses the GeneralName construct. If the email address
| includes an IDN but the local-part of the email address can be
| represented in ASCII, then the email address is placed in the
| rfc822Name choice of GeneralName, which is defined as type
| IA5String. If the local-part of the internationalized email
| address cannot be represented in ASCII, then the internationalized
| email address is placed in the otherName choice of GeneralName
| using the conventions in RFC 8398 [RFC8398].
|
| When the host-part contains an IDN, conforming implementations
| MUST convert all U-labels to A-labels.
|
| 7.5.1. Local-Part Contains Only ASCII Characters
|
| Two email addresses are considered to match if:
|
| 1) The local-part of each name is an exact match, AND
|
| 2) The host-part of each name matches using a case-insensitive
| ASCII comparison.
|
| Implementations that have a user interface SHOULD convert the
| host-part of internationalized email addresses specified in these
| extensions to Unicode before display. Specifically, conforming
| implementations convert A-labels to U-labels for display purposes.
|
| 7.5.2. Local-Part Contains Non-ASCII Characters
|
| When the local-part contains non-ASCII characters, conforming
| implementations MUST place the internationalized email address in
| the SmtpUTF8Mailbox within the otherName choice of GeneralName as
| specified in Section 3 of RFC 8398 [RFC8398]. Note that the UTF8
| encoding of the internationalized email address MUST NOT contain a
| Byte-Order-Mark (BOM) [RFC3629] to aid comparison. The email
| address local-part within the SmtpUTF8Mailbox MUST conform to the
| requirements of [RFC6530] and [RFC6531].
|
| Two email addresses are considered to match if:
|
| 1) The local-part of each name is an exact match, AND
|
| 2) The host-part of each name matches using a case-insensitive
| ASCII comparison.
|
| Implementations that have a user interface SHOULD convert the
| host-part of internationalized email addresses specified in these

| extensions to Unicode before display. Specifically, conforming
| implementations convert A-labels to U-labels for display purposes.

3. Security Considerations

The Security Considerations related to internationalized names in Section 4 of [RFC5890] are relevant to this specification.

Conforming Certification Authorities (CAs) SHOULD ensure that IDNs are valid according to IDNA2008, which is defined in [RFC5890], [RFC5891], [RFC5892], [RFC5893], [RFC5894], and the updates to these documents. Failure to use valid A-labels may yield a domain name that cannot be correctly represented in the Domain Name System (DNS). In addition, the CA/Browser Forum offers some guidance regarding internal server names in certificates [CABF].

An earlier version of this specification [RFC8399] required conversion of A-labels to U-labels in order to process name constraints for internationalized email addresses in SmtpUTF8Mailbox other names. This led to implementation complexity and at least two security vulnerabilities. Now, all IDNs are carried and processed as A-labels.

4. IANA Considerations

This document has no IANA actions.

5. References

5.1. 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, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3492] Costello, A., "Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)", RFC 3492, DOI 10.17487/RFC3492, March 2003, <https://www.rfc-editor.org/info/rfc3492>.

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

[RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, DOI 10.17487/RFC3987, January 2005, <https://www.rfc-editor.org/info/rfc3987>.

[RFC4518] Zeilenga, K., "Lightweight Directory Access Protocol (LDAP): Internationalized String Preparation", RFC 4518, DOI 10.17487/RFC4518, June 2006, <https://www.rfc-editor.org/info/rfc4518>.

[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, <https://www.rfc-editor.org/info/rfc5280>.

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

[RFC5891] Klensin, J., "Internationalized Domain Names in

Applications (IDNA): Protocol", RFC 5891,
DOI 10.17487/RFC5891, August 2010,
<https://www.rfc-editor.org/info/rfc5891>.

[RFC5892] Faltstrom, P., Ed., "The Unicode Code Points and Internationalized Domain Names for Applications (IDNA)", RFC 5892, DOI 10.17487/RFC5892, August 2010, <https://www.rfc-editor.org/info/rfc5892>.

[RFC5893] Alvestrand, H., Ed. and C. Karp, "Right-to-Left Scripts for Internationalized Domain Names for Applications (IDNA)", RFC 5893, DOI 10.17487/RFC5893, August 2010, <https://www.rfc-editor.org/info/rfc5893>.

[RFC6530] Klensin, J. and Y. Ko, "Overview and Framework for Internationalized Email", RFC 6530, DOI 10.17487/RFC6530, February 2012, <https://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, <https://www.rfc-editor.org/info/rfc6531>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8398] Melnikov, A., Ed. and W. Chuang, Ed., "Internationalized Email Addresses in X.509 Certificates", RFC 8398, DOI 10.17487/RFC8398, May 2018, <https://www.rfc-editor.org/info/rfc8398>.

5.2. Informative References

[CABF] CA/Browser Forum, "Internal Server Names and IP Address Requirements for SSL: Guidance on the Deprecation of Internal Server Names and Reserved IP Addresses provided by the CA/Browser Forum", Version 1.0, June 2012, <https://cabforum.org/internal-names/>.

[DDHQ] Datadog Security Labs, "The OpenSSL punycode vulnerability (CVE-2022-3602): Overview, detection, exploitation, and remediation", 1 November 2022, <https://securitylabs.datadoghq.com/articles/openssl-november-1-vulnerabilities/>.

[RFC3490] Faltstrom, P., Hoffman, P., and A. Costello, "Internationalizing Domain Names in Applications (IDNA)", RFC 3490, DOI 10.17487/RFC3490, March 2003, <https://www.rfc-editor.org/info/rfc3490>.

[RFC5894] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Background, Explanation, and Rationale", RFC 5894, DOI 10.17487/RFC5894, August 2010, <https://www.rfc-editor.org/info/rfc5894>.

[RFC8399] Housley, R., "Internationalization Updates to RFC 5280", RFC 8399, DOI 10.17487/RFC8399, May 2018, <https://www.rfc-editor.org/info/rfc8399>.

[UTS46] Davis, M. and M. Suignard, "Unicode Technical Standard #46: Unicode IDNA Compatibility Processing", Revision 31, The Unicode Consortium, Mountain View, September 2023, <https://www.unicode.org/reports/tr46>.

Acknowledgements

Thanks to David Benjamin and Wei Chuang for identifying the issue and a solution.

Thanks to Takahiro Nemoto, John Klensin, Mike Ounsworth, and Orie Steele for their careful review and thoughtful comments.

Author’s Address

Russ Housley
Vigil Security, LLC
Herndon, VA
United States of America
Email: housley@vigilsec.com