

Localized labels for testing IDNs

Cary Karp

Swedish Museum of Natural History

Abstract: Significant attention is currently being focused on the technical and policy bases for the introduction of top-level domains with labels derived from character strings with non-ASCII components, as well as on the growing inclusion of such labels in lower-level registries. Technical tests are being conducted in numerous contexts and this draft is intended to facilitate the comparability of their results. It describes a procedure for generating Punycode strings of arbitrary length that decode to typographically plausible sequences of Unicode characters in any desired script, without requiring particular understanding of unfamiliar writing systems. These strings are intended for testing the response of software applications to encoded sequences of ASCII characters of all possible lengths on all levels of a domain name, and the appearance of the corresponding Unicode strings in the display space. Constraints on the ability to base localized labels on established orthographies are also discussed.

Test requirements

The minimum length of a Punycode string is seven characters, encoding a single non-ASCII character. If the current policy constraints on one- and two-character labels are taken to apply to the displayed form, the minimum length of a stored Punycode string will be nine characters, corresponding to three non-ASCII characters. It should be noted, however, that these length restrictions were established when there was no difference between the stored and displayed forms of a label, and that some registries regard them as attaching to the stored form. It may be expected, nonetheless, that an IDN-aware revision of general policies can place the restriction on the number of displayed characters.

The longest TLD labels currently resolving in the root zone are the six-character .museum and .travel. These are both stored and displayed as ASCII characters but problems have been observed with the response of certain applications to them. However, since similar difficulty has been noted with shorter TLD labels, the problems are more likely to result from failure to recognize the labels as valid TLD designations, than simply from the length of the string. Assuming that requests will be made for new TLDs labeled with dictionary words with display lengths roughly equivalent to .museum and .travel, but written with non-ASCII characters, stored labels of twelve characters and longer will not be uncommon. If, as may also be expected, such things as the names of countries appear in full native representations, the lengths of the stored strings may be significantly greater. This suggests that stored strings of up to the maximum permitted length of 63 characters require evaluation. Even if there is no reason to expect that DNS resolvers will be taxed by the appearance of TLD labels of extreme length, one of the purposes of the technical testing is to identify unanticipated frailty. The response of other widely-deployed applications requires testing, in any case.

In many scripts, the way a character is displayed depends both on its position in a string and on the specific characters adjacent to it. If these shaping properties are to be manifested in the test environment as they are likely to appear in actual registered names, a test string cannot simply be a sequence of randomly selected characters. It can, however, be derived from a word taken from a dictionary of a language written with that script. If an online dictionary is available, its use will ease the determination of the requisite Unicode codepoints and avoid need for the manual transcription of unfamiliar scripts.

A distinction is made between conditions that pertain to laboratory testing in non-public namespaces, and those that attach to tests conducted in the public namespace. The former is focused on the generation of Punycode strings of varying length, with all display-side considerations relating solely to script. The latter, however, will involve linguistic considerations. Any label that is entered into the root zone of the DNS for the purposes of IDN testing will be categorically barred from subsequent delegation as a production domain. It is therefore also advisable to use a test term that would be unlikely to appear in the latter context, or is otherwise already restricted from such use. To reduce the potential for difficulty to an absolute minimum, a single word is therefore being recommended for all comparable test purposes.

The approach to generating test strings for use in private namespaces will be illustrated by deriving a non-lexical sequence from the word “hippopotamus” (a term likely to be found both in bilingual desk dictionaries and in corresponding online resources). To obviate any conceivable residual concern about rendering it inviable for subsequent candidacy for encoding into a production TLD label, a numerical sequence will be embedded in it. This is taken from the abbreviation “i18n” for “internationalization”, and is based (but not dependent) on the assumption that no TLD labels will have numerical components. No harm is likely to be done if the resulting string is unrecognizable in the language from which it was derived. The purpose is to generate typographically plausible sequences of characters in a variety of scripts, with no further semantic value or linguistic correctness being necessary (or even desirable). A purely lexical alternative will also be described for application in the public namespace.

Generating test labels

The effect of a TLD label that decodes to a string consisting of a single non-ASCII character (the minimum case) can be determined with the 17-character,

xn--flod18hst-12a

which is generated by the Swedish word *flodhäst* (a vernacular designation for the nearly extinct Swedish Forest Hippopotamus, as well as the non-indigenous hippopotamus species), with the interposed two digits,

flod18häst

This can simply be repeated to yield a Punycode string of arbitrary length, for example, the 27-character,

xn—flod18hstflod18hst-rtbj

from

flod18hästflod18häst

It will also be necessary to test ASCII-only sequences to distinguish between problems resulting solely from string length, and those caused by encoded non-ASCII characters, or otherwise by the Punycode prefix. One example of an ASCII-only string is the 56-character:

hippo18potamushippo18potamushippo18potamushippo18potamus

Total departure from the ASCII realm, using a simple alphabetic script written right-to-left, can be illustrated with the 28-character:

xn--18-rjdb cud0neb9a8ce1ezef

taken from the Yiddish:

חפאטאן18פאפא

This has the further advantage of testing embedded characters with opposing directional properties.

The properties of an alphabetic script with sophisticated shaping properties can be demonstrated with the 20-character,

xn--18-dtd1bdi0h3ask

from the Arabic,

فرس18النهر

No general assumption is made about scripts having syllable boundaries or other similarly convenient points for inserting digits into the test string, nor are any assumptions made about appropriate ways to proceed with non-alphabetic scripts. If the numerical device is not applicable, some other means should be applied to make the test sequence useless in the production environment. Such demonstration strings may be devised as appropriate for each desired script. (Note, however, that the string preparation process can place more or less severe constraints on the use of some scripts, as discussed below.)

One example of a label derived from a non-alphabetic script is the 15-character,

xn--18-h31ew85n

from the Kanji,

河18馬

Alternate approach

The procedure described above can be used to generate localized TLD labels for testing in the public namespace. However, that environment differs from the closed laboratory situation in one extremely significant regard – any recognizable linguistic attributes possessed by a string that resolves globally must be rigorously controlled to avoid its being seen as inappropriate by any corresponding speech community. In situations where that cannot be conveniently assured (or where the preceding approach is unnecessarily intricate), one option would be the straightforward use of benign dictionary terms.

A convenient vocabulary is provided by the RFC 2606 list of “Reserved Top Level DNS Names”, which explicitly lists four words that are restricted from autonomous delegation because:

“There is a need for top level domain (TLD) names that can be used for creating names which, without fear of conflicts with current or future actual TLD names in the global DNS, can be used for private testing of existing DNS related code, examples in documentation, DNS related experimentation, invalid DNS names, or other similar uses.”

Of the four names then reserved,

“‘.test’ is recommended for use in testing of current or new DNS related code”

and

“‘.example’ is recommended for use in documentation or as examples.”

The name “example” is also reserved on the second level in the .com, .net, and .org TLDs.

It would be counter to the conditions of RFC 2606 for either “.example” or “.test” to resolve in the root, but no restrictions are placed on lexical equivalents to those terms in other languages. One obvious alternative would therefore be to generate test TLD labels from translations of the word “test” into at least one language using each of the scripts that are represented in the public test. The second-level label in each such TLD could be similarly generated from the word “example”.

As noted above, any test label that is placed in the root zone will be unavailable for subsequent delegation. However, since the words “example” and “test” are already unavailable, similarly barring the equivalent words in any of the languages figuring in the public test (or perhaps generally in anticipation of future tests) would impose the smallest possible constraint on the production vocabulary. The translated “example” and “test” equivalents can therefore be used freely as IDN test strings in any situation where the intention is for them to be proper dictionary words. All requisite terms will also appear in any bilingual dictionary as discussed above. Examples of strings determined in this manner are उदाहरण.परीक्षा, 실례.테스트, and παράδειγμα.δοκιμή. If longer TLD labels are needed than those generated by a single instance of a translation of “test”, the word can be repeated as required, for example as, тест-тест-тест.

Finally, the utility of the publicly-accessible IDN test TLDs can easily extend beyond the technical trials specifically intended to precede the introduction of Punycode strings into the production root. It may therefore be worth keeping selected “.test” equivalents in persistent use. The utility of any such domain could be further enhanced by permitting the inclusion of second-level domains in addition to the “example.test” equivalents, as specific contributions to the development of IDN are put forward that would clearly benefit from appearing in this manner in the public namespace.

This demonstration registry might reasonably be placed under the administration of the IANA, with an appropriate reference group. It would be useful, in any case, to include a shareable list of all test strings generated from “example”, “hippopotamus” and “test”, plus any additional terms used for the purposes described here, in the “IANA Repository of TLD IDN Practices” at,

<http://www.iana.org/assignments/idn/>

General constraints

For differing reasons, all of the elements of a given writing system that might reasonably be requested for inclusion in a localized domain name may not actually be available. Some restrictions are inevitable consequences of the domain namespace never having been intended to serve as a vehicle for literary expression. Other limitations result from unanticipated problems and may be eliminated through protocol or policy revision. Converse difficulties have also resulted from excessive latitude in the available repertoire, and some currently viable characters may become unavailable both as IDNA is refined, and as registries adopt more restrictive policies.

Work is in underway in several contexts that is intended to clarify and rectify these issues. On first consideration, this might appear to be on a level of detail that is irrelevant to technical trials of the type described above. The agencies conducting such activity may, nonetheless, become engaged in dialog with local communities about specific constraints placed on their languages. This will require some familiarity with the kinds of limitations that are still being addressed, and the ability to assess the degree of transiency of specific issues. Key pending details are therefore reviewed below.

prohibits the correct representation of many names, such as that of the country Cyprus, which can only be incorrectly represented in IDN as “ΚΥΠΡΟΣ”. The German Eszett “ß” is similarly irrecoverably normalized to “ss” in the encoding process.

Appendix 1

The discussion of earlier versions of this text has revealed some misunderstanding about the first phase of the test action that has been commissioned in the interim. This is focused almost exclusively on the response of applications to Punycode strings of varying length when they appear in the top-level slot. The test strings were encoded from scripts selected on the basis of technical detail (as discussed in the body of this text). This selection was made without regard to demographic or other socio-linguistic factors that pertain to any languages that are written with these scripts.

Although the dictionaries of specific languages were consulted at the outset of the preparation of these test strings, the subsequent process methodically transformed the initial terms into character sequences that are not words. However, since the display form of a test string needs to have plausible typographic properties, some vestigial word-like appearance may remain. If any such sequence is discovered to be offensive or otherwise clearly inappropriate in an unforeseen cultural regard, it will simply be removed from the test roster.

Discussion of the prioritization of scripts as they appear in the public namespace in a later phase of the test needs to be conducted separately, and may reasonably apply to the instantiation of the fully lexical alternate approach described above. That action will be based on policy considerations that are still in early stages of consideration. The development of those policies will, in turn, be informed by the outcome of the impending technical trials in private namespaces.

Appendix 2

An extensive list of Punycode strings, including those listed above and extended with additional sequences needed for laboratory testing, will be appended to a subsequent version of this document.

Acknowledgements

This draft originated in a request by the ICANN President's Advisory Committee on IDN for a list of localized TLD labels for use in a series of impending technical trials. The text was then discussed by the working group that maintains the ICANN Guidelines for the Implementation of IDN, during the course of its work on rendering that document applicable to the top-level of the DNS. Further valuable commentary was received from an informal group of IETF members considering the need for modification of the IDNA protocol, and other colleagues.

All members of these groups are thanked for their advice, with particular acknowledgment of the contributions of the following people (in alphabetical order):

Harald Alvestrand
Tina Dam
Mohammed El Bashir
Michael Everson
Patrik Fältström
Hiro Hotta
Pat Kane
John Klensin
Ram Mohan

Author's address:

Cary Karp <ck@nrm.museum>
Swedish Museum of Natural History
Frescativägen 40
SE-10405 Stockholm