

Internet Engineering Task Force (IETF)
Request for Comments: 6497
Category: Informational
ISSN: 2070-1721

M. Davis
Google
A. Phillips
Lab126
Y. Umaoka
IBM
C. Falk
Infinite Automata
February 2012

BCP 47 Extension T - Transformed Content

Abstract

This document specifies an Extension to BCP 47 that provides subtags for specifying the source language or script of transformed content, including content that has been transliterated, transcribed, or translated, or in some other way influenced by the source. It also provides for additional information used for identification.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6497>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	4
2. BCP 47 Required Information	4
2.1. Overview	4
2.2. Structure	6
2.3. Canonicalization	7
2.4. BCP 47 Registration Form	8
2.5. Field Definitions	8
2.6. Registration of Field Subtags	10
2.7. Registration of Additional Fields	11
2.8. Committee Responses to Registration Proposals	11
2.9. Machine-Readable Data	11
3. Acknowledgements	14
4. IANA Considerations	14
5. Security Considerations	14
6. References	14
6.1. Normative References	14
6.2. Informative References	15

1. Introduction

[BCP47] permits the definition and registration of language tag extensions "that contain a language component and are compatible with applications that understand language tags". This document defines an extension for specifying the source of content that has been transformed, including text that has been transliterated, transcribed, or translated, or in some other way influenced by the source. It may be used in queries to request content that has been transformed. The "singleton" identifier for this extension is 't'.

Language tags, as defined by [BCP47], are useful for identifying the language of content. There are mechanisms for specifying variant subtags for special purposes. However, these variants are insufficient for specifying content that has undergone transformations, including content that has been transliterated, transcribed, or translated. The correct interpretation of the content may depend upon knowledge of the conventions used for the transformation.

Suppose that Italian or Russian cities on a map are transcribed for Japanese users. Each name needs to be transliterated into katakana using rules appropriate for the specific source and target language. When tagging such data, it is important to be able to indicate not only the resulting content language ("ja" in this case), but also the source language.

Transforms such as transliterations may vary, depending not only on the basis of the source and target script, but also on the source and target language. Thus, the Russian <U+041F U+0443 U+0442 U+0438 U+043D> (which corresponds to the Cyrillic <PE, U, TE, I, EN>) transliterates into "Putin" in English but "Poutine" in French. The identifier could be used to indicate a desired mechanical transformation in an API, or could be used to tag data that has been converted (mechanically or by hand) according to a transliteration method.

In addition, many different conventions have arisen for how to transform text, even between the same languages and scripts. For example, "Gaddafi" is commonly transliterated from Arabic to English as any of (G/Q/K/Kh)a(d/dh/dd/dhdh/th/zz)af(i/y). Some examples of standardized conventions used for transcribing or transliterating text include:

- a. United Nations Group of Experts on Geographical Names (UNGEGN)
- b. US Library of Congress (LOC)
- c. US Board on Geographic Names (BGN)
- d. Korean Ministry of Culture, Sports and Tourism (MCST)
- e. International Organization for Standardization (ISO)

The usage of this extension is not limited to formal transformations, and may include other instances where the content is in some other way influenced by the source. For example, this extension could be used to designate a request for a speech recognizer that is tailored

specifically for second-language speakers who are first-language speakers of a particular language (e.g., a recognizer for "English spoken with a Chinese accent").

1.1. 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].

2. BCP 47 Required Information

2.1. Overview

Identification of transformed content can be done using the 't' extension defined in this document. This extension is formed by the 't' singleton followed by a sequence of subtags that would form a language tag as defined by [BCP47]. This allows the source language or script to be specified to the degree of precision required. There are restrictions on the sequence of subtags. They MUST form a regular, valid, canonical language tag, and MUST neither include extensions nor private use sequences introduced by the singleton 'x'. Where only the script is relevant (such as identifying a script-script transliteration), then 'und' is used for the primary language subtag.

For example:

Language Tag	Description
ja-t-it	The content is Japanese, transformed from Italian.
ja-Kana-t-it	The content is Japanese Katakana, transformed from Italian.
und-Latn-t-und-cyrl	The content is in the Latin script, transformed from the Cyrillic script.

Note that the sequence of subtags governed by 't' cannot contain a singleton (a single-character subtag), because that would start a new extension. For example, the tag "ja-t-i-ami" does not indicate that the source is in "i-ami", because "i-ami" is not a regular language tag in [BCP47]. That tag would express an empty 't' extension followed by an 'i' extension.

The 't' extension is not intended for use in structured data that already provides separate source and target language identifiers. For example, this is the case in localization interchange formats such as XLIFF. In such cases, it would be inappropriate to use "ja-t-it" for the target language tag because the source language tag "it" would already be present in the data. Instead, one would use the language tag "ja".

As noted earlier, it is sometimes necessary to indicate additional information about a transformation. This additional information is optionally supplied after the source in a series of one or more fields, where each field consists of a field separator subtag followed by one or more non-separator subtags. Each field separator subtag consists of a single letter followed by a single digit.

A transformation mechanism is an optional field that indicates the specification used for the transformation, such as "UNGEGN" for the United Nations Group of Experts on Geographical Names transliterations and transcriptions. It uses the 'm0' field separator followed by certain subtags.

For example:

Language Tag	Description
und-Cyrl-t-und-latn-m0-ungegn-2007	The content is in Cyrillic, transformed from Latin, according to a UNGEGN specification dated 2007.

The field separator subtags, such as 'm0', were chosen because they are short, visually distinctive, and cannot occur in a language subtag (outside of an extension and after 'x'), thus eliminating the potential for collision or confusion with the source language tag.

The field subtags are defined by Section 3 of Unicode Technical Standard #35: Unicode Locale Data Markup Language (LDML) [UTS35], the main specification for the Unicode Common Locale Data Repository (CLDR) project. That section also defines the parallel 'u' extension [RFC6067], for which the Unicode Consortium is also the maintaining authority. As required by BCP 47, subtags follow the language tag ABNF and other rules for the formation of language tags and subtags, are restricted to the ASCII letters and digits, are not case sensitive, and do not exceed eight characters in length.

The LDML specification is available over the Internet and at no cost, and is available via a royalty-free license at <http://unicode.org/copyright.html>. LDML is versioned, and each version of LDML is numbered, dated, and stable. Extension subtags, once defined by LDML, are never retracted or substantially changed in meaning.

The maintaining authority for the 't' extension is the Unicode Consortium:

Item	Value
Name	Unicode Consortium
Contact Email	cldr-contact@unicode.org
Discussion	cldr-users@unicode.org
List Email	
URL Location	cldr.unicode.org
Specification	Unicode Technical Standard #35 Unicode Locale Data Markup Language (LDML), http://unicode.org/reports/tr35/
Section	Section 3 Unicode Language and Locale Identifiers

2.2. Structure

The subtags in the 't' extension are of the following form:

```
t-ext      = "t"                ; Extension
            (( "-" lang *("-" field)) ; Source + optional field(s)
            / 1*("-" field))       ; Field(s) only (no source)

lang       = language          ; BCP 47, with restrictions
            [ "-" script]
            [ "-" region]
            *("-" variant)

field      = fsep 1*("-" 3*8alphanum) ; With restrictions

fsep      = ALPHA DIGIT          ; Subtag separators
alphanum  = ALPHA / DIGIT
```

where <language>, <script>, <region>, and <variant> rules are specified in [BCP47], and <ALPHA> and <DIGIT> rules in [RFC5234].

Description and restrictions:

- a. The 't' extension MUST have at least one subtag.
- b. The 't' extension normally starts with a source language tag, which MUST be a regular, canonical language tag as specified by [BCP47]. Tags described by the 'irregular' production in BCP 47 MUST NOT be used to form the language tag. The source language tag MAY be omitted: some field values do not require it.
- c. There is optionally a sequence of fields, where each field has a separator followed by a sequence of one or more subtags. Two identical field separators MUST NOT be present in the language tag.
- d. The order of the fields in a 't' extension is not significant. The order of subtags within a field is significant. See Section 2.3 ("Canonicalization").
- e. The 't' subtag fields are defined by Section 3 of Unicode Technical Standard #35: Unicode Locale Data Markup Language [UTS35].

2.3. Canonicalization

As required by [BCP47], the use of uppercase or lowercase letters is not significant in the subtags used in this extension. The canonical form for all subtags in the extension is lowercase, with the fields ordered by the separators, alphabetically. The order of subtags within a field is significant, and MUST NOT be changed in the process of canonicalizing.

2.4. BCP 47 Registration Form

Per RFC 5646, Section 3.7 [BCP47]:

%%

Identifier: t

Description: Specifying Transformed Content

Comments: Subtags for the identification of content that has been transformed, including but not limited to:
transliteration, transcription, and translation.

Added: 2011-12-16

RFC: RFC 6497

Authority: Unicode Consortium

Contact_Email: cldr-contact@unicode.org

Mailing_List: cldr-users@unicode.org

URL: <http://www.unicode.org/Public/cldr/latest/core.zip>

%%

2.5. Field Definitions

Assignment of 't' field subtags is determined by the Unicode CLDR Technical Committee, in accordance with the policies and procedures in <http://www.unicode.org/consortium/tc-procedures.html>, and subject to the Unicode Consortium Policies on <http://www.unicode.org/policies/policies.html>.

Assignments that can be made by successive versions of LDML [UTS35] by the Unicode Consortium without requiring a new RFC include:

- o The allocation of new field separator subtags for use after the 't' extension.
- o The allocation of subtags valid after a field separator subtag.
- o The addition of subtag aliases and descriptions.
- o The modification of subtag descriptions.

Changes to the syntax or meaning of the 't' extension would require a new RFC that obsoletes this document; such an RFC would break stability, and would thus be contrary to the policies of the Unicode Consortium.

At the time this document was published, one field separator subtag was specified in [UTS35]: the transform mechanism. That field is summarized here:

- a. The transform mechanism consists of a sequence of subtags starting with the 'm0' separator followed by one or more mechanism subtags. Each mechanism subtag has a length of 3 to 8 alphanumeric characters. The sequence as a whole provides an identification of the specification for the transform, such as the mechanism subtag 'ungegn' in "und-Cyrl-t-und-latn-m0-ungegn". In many cases, only one mechanism subtag is necessary, but multiple subtags MAY be defined in [UTS35] where necessary.
- b. Any purely numeric subtag is a representation of a date in the Gregorian calendar. It MAY occur in any mechanism field, but it SHOULD only be used where necessary. If it does occur:
 - * it MUST occur as the final subtag in the field
 - * it MUST NOT be the only subtag in the field
 - * it MUST only consist of a sequence of digits of the form YYYY, YYYYMM, or YYYYMMDD
 - * it SHOULD be as short as possible

Note: The format is related to that of [RFC3339], but is not the same. The RFC 3339 full-date won't work because it uses hyphens. The offset ("Z") is not used because the date is a publication date (aka 'floating date'). For more information, see Section 3.3 ("Floating Time") of [W3C-TimeZones].

c. Examples:

- * 20110623 represents June 23, 2011.
- * There are three dated versions of the UNGEGN transliteration specification for Hebrew to Latin. They can be represented by the following language tags:
 - + und-Hebr-t-und-latn-m0-ungegn-1972
 - + und-Hebr-t-und-latn-m0-ungegn-1977
 - + und-Hebr-t-und-latn-m0-ungegn-2007

- * Suppose that the BGN transliteration specification for Cyrillic to Latin had three versions, dated June 11, 1999; Dec 30, 1999; and May 1, 2011. In that case, the corresponding first two DATE subtags would require the months to be distinctive (199906 and 199912), but the last subtag would only require the year (2011).

- d. Some mechanisms may use a versioning system that is not distinguished by date, or not by date alone. In the latter case, the version will be of a form specified by [UTS35] for that mechanism. For example, if the mechanism xxx uses versions of the form v21a, then a tag could look like "ja-t-it-m0-xxx-v21a". If there are multiple sub-versions distinguished by date, then a tag could look like "ja-t-it-m0-xxx-v21a-2007".

A language tag with the 't' extension MAY be used to request a specific transform of content. In such a case, the recipient SHOULD return content that corresponds as closely as feasible to the requested transform, including the specification of the mechanism. For example, if the request is ja-t-it-m0-xxx-v21a-2007, and the recipient has content corresponding to both ja-t-it-m0-xxx-v21a and ja-t-it-m0-xxx-v21b-2009, then the v21a version would be preferred. As is the case for language matching as discussed in [BCP47], different implementations MAY have different measures of "closeness".

2.6. Registration of Field Subtags

Registration of transform mechanisms is requested by filing a ticket at <http://cldr.unicode.org/>. The proposal in the ticket MUST contain the following information:

Item	Description
Subtag	The proposed mechanism subtag (or subtag sequence).
Description	A description of the proposed mechanism; that description MUST be sufficient to distinguish it from other mechanisms in use.
Version	If versioning for the mechanism is not done according to date, then a description of the versioning conventions used for the mechanism.

Proposals for clarifications of descriptions or additional aliases may also be requested by filing a ticket.

The committee MAY define a template for submissions that requests more information, if it is found that such information would be useful in evaluating proposals.

2.7. Registration of Additional Fields

In the event that it proves necessary to add an additional field (such as 'm2'), it can be requested by filing a ticket at <http://cldr.unicode.org/>. The proposal in the ticket MUST contain a full description of the proposed field semantics and subtag syntax, and MUST conform to the ABNF syntax for "field" presented in Section 2.2.

2.8. Committee Responses to Registration Proposals

The committee MUST post each proposal publicly within 2 weeks after reception, to allow for comments. The committee must respond publicly to each proposal within 4 weeks after reception.

The response MAY:

- o request more information or clarification
- o accept the proposal, optionally with modifications to the subtag or description
- o reject the proposal, because of significant objections raised on the mailing list or due to problems with constraints in this document or in [UTS35]

Accepted tickets result in a new entry in the machine-readable CLDR BCP 47 data or, in the case of a clarified description, modifications to the description attribute value for an existing entry.

2.9. Machine-Readable Data

Beginning with CLDR version 1.7.2, machine-readable files are available listing the data defined for BCP 47 extensions for each successive version of [UTS35]. The data in these files is used for testing the validity of subtags for the 't' extension and for the 'u' extension [RFC6067], for which the Unicode Consortium is also the maintaining authority. These releases are listed on <http://cldr.unicode.org/index/downloads>. Each release has an associated data directory of the form "http://unicode.org/Public/cldr/<version>", where "<version>" is replaced by the release number. For example, for version 1.7.2, the

"core.zip" file is located at <http://unicode.org/Public/cldr/1.7.2/core.zip>. The most recent version is always identified by the version "latest" and can be accessed by the URL in Section 2.4.

Inside the "core.zip" file, the directory "common/bcp47" contains the data files listing the valid attributes, keys, and types for each successive version of [UTS35]. Each data file lists the keys and types relevant to that topic.

The XML structure lists the keys, such as `<key extension="t" name="m0" description="Transliteration extension mechanism">`, with subelements for the types, such as `<type name="ungegn" description="United Nations Group of Experts on Geographical Names"/>`. The currently defined attributes for the mechanisms include:

Attribute	Description	Examples
name	The name of the mechanism, limited to 3-8 characters (or sequences of them).	UNGEGN, ALALC
description	A description of the name, with all and only that information necessary to distinguish one name from others with which it might be confused. Descriptions are not intended to provide general background information.	United Nations Group of Experts on Geographical Names; American Library Association-Library of Congress
since	Indicates the first version of CLDR where the name appears. (Required for new items.)	1.9, 2.0.1
alias	Alternative name of the key or type, not limited in number of characters. Aliases are intended for backwards compatibility, not to provide all possible alternate names or designations. (Optional.)	

The file for the transform extension is "transform.xml". The initial version of that file contains the following information.

```
<keyword>
  <key extension="t" name="m0" description=
    "Transliteration extension mechanism">
    <type name="ungegn" description=
      "United Nations Group of Experts on Geographical Names"
      since="21"/>
    <type name="alaloc" description=
      "American Library Association-Library of Congress"
      since="21"/>
    <type name="bgn" description=
      "US Board on Geographic Names"
      since="21"/>
    <type name="mcst" description=
      "Korean Ministry of Culture, Sports and Tourism"
      since="21"/>
    <type name="iso" description=
      "International Organization for Standardization"
      since="21"/>
    <type name="din" description=
      "Deutsches Institut fuer Normung"
      since="21"/>
    <type name="gost" description=
      "Euro-Asian Council for Standardization, Metrology
      and Certification"
      since="21"/>
  </key>
</keyword>
```

To get the version information in XML when working with the data files, the XML parser must be validating. When the 'core.zip' file is unzipped, the 'dtd' directory will be at the same level as the 'bcp47' directory; this is required for correct validation. For each release after CLDR 1.8, types introduced in that release are also marked in the data files by the XML attribute "since", such as in the following example:

```
<type name="adp" since="1.9"/>
```

The data is also currently maintained in a source code repository, with each release tagged, for viewing directly without unzipping. For example, see:

- o <http://unicode.org/repos/cldr/tags/release-1-7-2/common/bcp47/>
- o <http://unicode.org/repos/cldr/tags/release-1-8/common/bcp47/>

For more information, see
<http://cldr.unicode.org/index/bcp47-extension>.

3. Acknowledgements

Thanks to John Emmons and the rest of the Unicode CLDR Technical Committee for their work in developing the BCP 47 subtags for LDML.

4. IANA Considerations

IANA has inserted the record of Section 2.4 into the Language Extensions Registry, according to Section 3.7 ("Extensions and the Extensions Registry") of "Tags for Identifying Languages" [BCP47]. Per Section 5.2 of [BCP47], there might be occasional (rare) requests by the Unicode Consortium (the "Authority" listed in the record) for maintenance of this record. Changes that can be submitted to IANA without the publication of a new RFC are limited to modification of the Comments, Contact_Email, Mailing_List, and URL fields. Any such requested changes MUST use the domain 'unicode.org' in any new addresses or URIs, MUST explicitly cite this document (so that IANA can reference these requirements), and MUST originate from the 'unicode.org' domain. The domain or authority can only be changed via a new RFC.

5. Security Considerations

The security considerations for this extension are the same as those for [BCP47]. See RFC 5646, Section 6, Security Considerations [BCP47].

6. References

6.1. Normative References

- [BCP47] Phillips, A., Ed., and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, September 2009.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5234] Crocker, D., Ed., and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [UTS35] Davis, M., "Unicode Technical Standard #35: Locale Data Markup Language (LDML)", February 2012, <<http://www.unicode.org/reports/tr35/>>.

6.2. Informative References

- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, July 2002.
- [RFC6067] Davis, M., Phillips, A., and Y. Umaoka, "BCP 47 Extension U", RFC 6067, December 2010.
- [W3C-TimeZones]
Phillips, Ed., "W3C Working Group Note: Working with Time Zones", July 2011,
<<http://www.w3.org/TR/2011/NOTE-timezone-20110705/>>.

Authors' Addresses

Mark Davis
Google

E-Mail: mark@macchiato.com

Addison Phillips
Lab126

E-Mail: addison@lab126.com

Yoshito Umaoka
IBM

E-Mail: yoshito_umaoka@us.ibm.com

Courtney Falk
Infinite Automata

E-Mail: court@infiauto.com