

NAME

unicode - universal character set

DESCRIPTION

The international standard ISO 10646 defines the Universal Character Set (UCS). UCS contains all characters of all other character set standards. It also guarantees "round-trip compatibility"; in other words, conversion tables can be built such that no information is lost when a string is converted from any other encoding to UCS and back.

UCS contains the characters required to represent practically all known languages. This includes not only the Latin, Greek, Cyrillic, Hebrew, Arabic, Armenian, and Georgian scripts, but also Chinese, Japanese and Korean Han ideographs as well as scripts such as Hiragana, Katakana, Hangul, Devanagari, Bengali, Gurmukhi, Gujarati, Oriya, Tamil, Telugu, Kannada, Malayalam, Thai, Lao, Khmer, Bopomofo, Tibetan, Runic, Ethiopic, Canadian Syllabics, Cherokee, Mongolian, Ogham, Myanmar, Sinhala, Thaana, Yi, and others. For scripts not yet covered, research on how to best encode them for computer usage is still going on and they will be added eventually. This might eventually include not only Hieroglyphs and various historic Indo-European languages, but even some selected artistic scripts such as Tengwar, Cirth, and Klingon. UCS also covers a large number of graphical, typographical, mathematical, and scientific symbols, including those provided by TeX, Postscript, APL, MS-DOS, MS-Windows, Macintosh, OCR fonts, as well as many word processing and publishing systems, and more are being added.

The UCS standard (ISO 10646) describes a 31-bit character set architecture consisting of 128 24-bit *groups*, each divided into 256 16-bit *planes* made up of 256 8-bit *rows* with 256 *column* positions, one for each character. Part 1 of the standard (ISO 10646-1) defines the first 65534 code positions (0x0000 to 0xffff), which form the *Basic Multilingual Plane* (BMP), that is plane 0 in group 0. Part 2 of the standard (ISO 10646-2) adds characters to group 0 outside the BMP in several *supplementary planes* in the range 0x10000 to 0x10ffff. There are no plans to add characters beyond 0x10ffff to the standard, therefore of the entire code space, only a small fraction of group 0 will ever be actually used in the foreseeable future. The BMP contains all characters found in the commonly used other character sets. The supplemental planes added by ISO 10646-2 cover only more exotic characters for special scientific, dictionary printing, publishing industry, higher-level protocol and enthusiast needs.

The representation of each UCS character as a 2-byte word is referred to as the UCS-2 form (only for BMP characters), whereas UCS-4 is the representation of each character by a 4-byte word. In addition, there exist two encoding forms UTF-8 for backward compatibility with ASCII processing software and UTF-16 for the backward-compatible handling of non-BMP characters up to 0x10ffff by UCS-2 software.

The UCS characters 0x0000 to 0x007f are identical to those of the classic US-ASCII character set and the characters in the range 0x0000 to 0x00ff are identical to those in ISO 8859-1 (Latin-1).

Combining characters

Some code points in UCS have been assigned to *combining characters*. These are similar to the nonspacing accent keys on a typewriter. A combining character just adds an accent to the previous character. The most important accented characters have codes of their own in UCS, however, the combining character mechanism allows us to add accents and other diacritical marks to any character. The combining characters always follow the character which they modify. For example, the German character Umlaut-A ("Latin capital letter A with diaeresis") can either be represented by the precomposed UCS code 0x00c4, or alternatively as the combination of a normal "Latin capital letter A" followed by a "combining diaeresis": 0x0041 0x0308.

Combining characters are essential for instance for encoding the Thai script or for mathematical typesetting and users of the International Phonetic Alphabet.

Implementation levels

As not all systems are expected to support advanced mechanisms like combining characters, ISO 10646-1 specifies the following three *implementation levels* of UCS:

Level 1 Combining characters and Hangul Jamo (a variant encoding of the Korean script, where a Hangul syllable glyph is coded as a triplet or pair of vowel/consonant codes) are not supported.

- Level 2 In addition to level 1, combining characters are now allowed for some languages where they are essential (e.g., Thai, Lao, Hebrew, Arabic, Devanagari, Malayalam).
- Level 3 All UCS characters are supported.

The Unicode 3.0 Standard published by the Unicode Consortium contains exactly the UCS Basic Multilingual Plane at implementation level 3, as described in ISO 10646-1:2000. Unicode 3.1 added the supplemental planes of ISO 10646-2. The Unicode standard and technical reports published by the Unicode Consortium provide much additional information on the semantics and recommended usages of various characters. They provide guidelines and algorithms for editing, sorting, comparing, normalizing, converting, and displaying Unicode strings.

Unicode under Linux

Under GNU/Linux, the C type *wchar_t* is a signed 32-bit integer type. Its values are always interpreted by the C library as UCS code values (in all locales), a convention that is signaled by the GNU C library to applications by defining the constant `__STDC_ISO_10646__` as specified in the ISO C99 standard.

UCS/Unicode can be used just like ASCII in input/output streams, terminal communication, plaintext files, filenames, and environment variables in the ASCII compatible UTF-8 multibyte encoding. To signal the use of UTF-8 as the character encoding to all applications, a suitable *locale* has to be selected via environment variables (e.g., "LANG=en_GB.UTF-8").

The `nl_langinfo(CODESET)` function returns the name of the selected encoding. Library functions such as `wctomb(3)` and `mbsrtowcs(3)` can be used to transform the internal *wchar_t* characters and strings into the system character encoding and back and `wcwidth(3)` tells, how many positions (0–2) the cursor is advanced by the output of a character.

Private Use Areas (PUA)

In the Basic Multilingual Plane, the range 0xe000 to 0xf8ff will never be assigned to any characters by the standard and is reserved for private usage. For the Linux community, this private area has been subdivided further into the range 0xe000 to 0xefff which can be used individually by any end-user and the Linux zone in the range 0xf000 to 0xf8ff where extensions are coordinated among all Linux users. The registry of the characters assigned to the Linux zone is maintained by LANANA and the registry itself is *Documentation/unicode.txt* in the Linux kernel sources.

Two other planes are reserved for private usage, plane 15 (Supplementary Private Use Area-A, range 0xf0000 to 0xffffd) and plane 16 (Supplementary Private Use Area-B, range 0x100000 to 0x10ffffd).

Literature

- * Information technology — Universal Multiple-Octet Coded Character Set (UCS) — Part 1: Architecture and Basic Multilingual Plane. International Standard ISO/IEC 10646-1, International Organization for Standardization, Geneva, 2000.

This is the official specification of UCS . Available from [Unknown](#).

- * The Unicode Standard, Version 3.0. The Unicode Consortium, Addison-Wesley, Reading, MA, 2000, ISBN 0-201-61633-5.
- * S. Harbison, G. Steele. C: A Reference Manual. Fourth edition, Prentice Hall, Englewood Cliffs, 1995, ISBN 0-13-326224-3.

A good reference book about the C programming language. The fourth edition covers the 1994 Amendment 1 to the ISO C90 standard, which adds a large number of new C library functions for handling wide and multibyte character encodings, but it does not yet cover ISO C99, which improved wide and multibyte character support even further.

- * Unicode Technical Reports.

[Unknown](#)

- * Markus Kuhn: UTF-8 and Unicode FAQ for UNIX/Linux.

[Unknown](#)

- * Bruno Haible: Unicode HOWTO.
[Unknown](#)

SEE ALSO

[locale\(1\)](#), [setlocale\(3\)](#), [charsets\(7\)](#), [utf-8\(7\)](#)

COLOPHON

This page is part of release 4.10 of the Linux *man-pages* project. A description of the project, information about reporting bugs, and the latest version of this page, can be found at <https://www.kernel.org/doc/man-pages/>.