NAME

tzfile - timezone information

DESCRIPTION

This page describes the structure of the timezone files used by tzset(3). These files are typically found under one of the directories */usr/lib/zoneinfo* or */usr/share/zoneinfo*.

Timezone information files begin with a 44-byte header structured as follows:

* The magic four-byte sequence "TZif" identifying this as a timezone information file.
* A single character identifying the version of the file's format: either an ASCII NUL ('\0') or a '2' (**0x32**).
* Fifteen bytes containing zeros reserved for future use.
* Six four-byte values of type *long*, written in a "standard" byte order (the high-order byte of the value is written first). These values are, in order:

  tzh_ttisgmtcnt
  : The number of UTC/local indicators stored in the file.

  tzh_ttisstdcnt
  : The number of standard/wall indicators stored in the file.

  tzh_leapcnt
  : The number of leap seconds for which data is stored in the file.

  tzh_timecnt
  : The number of "transition times" for which data is stored in the file.

  tzh_typecnt
  : The number of "local time types" for which data is stored in the file (must not be zero).

  tzh_charcnt
  : The number of characters of "timezone abbreviation strings" stored in the file.

The above header is followed by *tzh_timecnt* four-byte values of type *long*, sorted in ascending order. These values are written in "standard" byte order. Each is used as a transition time (as returned by time(2)) at which the rules for computing local time change. Next come *tzh_timecnt* one-byte values of type *unsigned char*; each one tells which of the different types of "local time" types described in the file is associated with the same-indexed transition time. These values serve as indices into an array of *ttinfo* structures (with *tzh_typecnt* entries) that appear next in the file; these structures are defined as follows:

```
struct ttinfo {
long tt_gmtoff;
int tt_isdst;
unsigned int tt_abbrind;
};
```

Each structure is written as a four-byte value for *tt_gmtoff* of type *long*, in a standard byte order, followed by a one-byte value for *tt_isdst* and a one-byte value for *tt_abbrind*. In each structure,*tt_gmtoff* gives the number of seconds to be added to UTC, *tt_isdst* tells whether *tm_isdst* should be set by localtime(3), and *tt_abbrind* serves as an index into the array of timezone abbreviation characters that follow the *ttinfo* structure(s) in the file.

Then there are *tzh_leapcnt* pairs of four-byte values, written in standard byte order; the first value of each pair gives the time (as returned by time(2)) at which a leap second occurs; the second gives the *total* number of leap seconds to be applied after the given time. The pairs of values are sorted in ascending order by time.

Then there are *tzh_ttisstdcnt* standard/wall indicators, each stored as a one-byte value; they tell whether the transition times associated with local time types were specified as standard time or wall clock time, and are used when a timezone file is used in handling POSIX-style timezone environment variables.

Finally, there are *tzh_ttisgmtcnt* UTC/local indicators, each stored as a one-byte value; they tell whether the

transition times associated with local time types were specified as UTC or local time, and are used when a timezone file is used in handling POSIX-style timezone environment variables.

localtime(3) uses the first standard-time *ttinfo* structure in the file (or simply the first *ttinfo* structure in the absence of a standard-time structure) if either *tzh_timecnt* is zero or the time argument is less than the first transition time recorded in the file.

NOTES

This manual page documents <*tzfile.h*> in the glibc source archive, see *timezone/tzfile.h*.

It seems that timezone uses **tzfile** internally, but glibc refuses to expose it to userspace. This is most likely because the standardised functions are more useful and portable, and actually documented by glibc. It may only be in glibc just to support the non-glibc-maintained timezone data (which is maintained by some other entity).

Version 2 format

For version-2-format timezone files, the above header and data is followed by a second header and data, identical in format except that eight bytes are used for each transition time or leap-second time (and that the version byte in the header record is **0x32** rather than **0x00**). After the second header and data comes a newline-enclosed, POSIX-TZ-environment-variable-style string for use in handling instants after the last transition time stored in the file (with nothing between the newlines if there is no POSIX representation for such instants).

The second section of the timezone file consists of another 44-byte header record, identical in structure to the one at the beginning of the file, except that it applies to the data that follows, which is also identical in structure to the first section of the timezone file, with the following differences:

* The transition time values, after the header, are eight-byte values.
* In each leap second record, the leap second value is an eight-byte value. The accumulated leap second count is still a four-byte value.

In all cases, the eight-byte time values are given in the "standard" byte order, the high-order byte first.

POSIX timezone string

The second eight-byte time value section is followed by an optional third section: a single ASCII newline character ('\n'), then a text string followed by a second newline character. The text string is a POSIX timezone string, whose format is described in the tzset(3) manual page.

The POSIX timezone string defines a rule for computing transition times that follow the last transition time explicitly specified in the timezone information file.

Summary of the timezone information file format

Four-byte value section
(header version **0x00** or **0x32**)
Header record
Four-byte transition times
Transition time index
ttinfo structures
Timezone abbreviation array
Leap second records
Standard/Wall array
UTC/Local array

Eight-byte value section
(only if first header version is **0x32**,
the second header's version is also **0x32**)
Header record
Eight-byte transition times
Transition time index

ttinfo structures
Timezone abbreviation array
Leap second records
Standard/Wall array
UTC/Local array

Third section
(optional, only in **0x32** version files)
Newline character
Timezone string
Newline character

SEE ALSO

ctime(3), tzset(3), tzselect(8),

timezone/tzfile.h in the glibc source tree

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/.