

Independent Submission
Internet-Draft
Updates: [7808](#) (if approved)
Intended status: Standards Track
Expires: December 7, 2018

A. Olson

P. Eggert
UCLA
K. Murchison
FastMail
June 5, 2018

The Time Zone Information Format (TZif)
draft-murchison-tzdist-tzif-05

Abstract

This document defines the Time Zone Information Format (TZif) for representing and exchanging time zone information, independent of any particular service or protocol. A MIME media type for this format is also defined.

Open Issues

- o Should we require TZDIST servers to always return a particular format (2 or 3) with/without leap seconds? Or should we allow the client to choose what they want (via parameters on the MIME type)? E.g. application/tzif+leap; version=3
- o Is the MIME registration complete and correct (e.g. any security or privacy considerations)?

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 7, 2018.

Copyright Notice

Copyright (c) 2018 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 (<https://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.

Table of Contents

1.	Introduction	2
2.	Conventions Used in This Document	3
3.	The Time Zone Information Format (TZif)	4
3.1.	TZif Header	5
3.2.	TZif Data Block	7
3.3.	TZif Footer	10
3.3.1.	TZ String Extensions	11
4.	Use with the Time Zone Data Distribution Service	11
4.1.	Example	11
5.	Security Considerations	13
6.	Privacy Considerations	13
7.	IANA Considerations	13
8.	Acknowledgments	14
9.	References	14
9.1.	Normative References	14
9.2.	Informative References	15
9.3.	URIs	15
Appendix A.	Change History (To be removed by RFC Editor before publication)	15
	Authors' Addresses	17

[1.](#) Introduction

Time zone data typically consists of offsets from Universal Time (UT), daylight saving transition rules, one or more local time designations (acronyms or abbreviations), and optional leap second adjustments. One such format for conveying this information is iCalendar [[RFC5545](#)]. It is a text-based format used by calendaring and scheduling systems.

This document defines the Time Zone Information Format. It is a binary format used by most UNIX systems to calculate local time. There is a wide variety of interoperable software [[tz-link](#)] capable of generating and reading files in this format.

This specification does not define the source of the time zone data or leap second information. One such source is the IANA-hosted time zone database [[RFC6557](#)].

2. Conventions Used in This Document

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](#) [[1](#)] [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

The following terms are used in this document:

Coordinated Universal Time (UTC): The basis for civil time since 1960. It is approximately equal to mean solar time at the prime meridian (0 degrees longitude).

Daylight Saving Time (DST): The time according to a location's law or practice, adjusted as necessary from standard time. The adjustment may be positive, negative, or zero.

International Atomic Time (TAI): The time standard based on atomic clocks. It is equal to UTC except without leap second adjustments.

Local Time: The time according to a location's current time zone offset from Universal Time.

Standard Time: The time according to a location's law or practice, unadjusted for Daylight Saving Time.

Time Change: A change to civil timekeeping practice. It occurs when one or more of the following happen simultaneously:

1. a change in UT offset
2. a change in whether standard or daylight saving time is in use
3. a change in time zone abbreviation
4. a leap second (i.e., a change in UTC - TAI)

Time Zone Data: The Time Zone Data Distribution Service (TZDIST) [[RFC7808](#)] defines "Time zone data" as "data that defines a single time zone, including an identifier, UT offset values, DST rules, and other information such as time zone abbreviations." The

interchange format defined in this document is one such form of time zone data.

Universal Time (UT): The basis of civil time. This is the principal form of the mean solar time at the prime meridian (0 degrees longitude) for timestamps before UTC was introduced in 1960, and is UTC for timestamps thereafter. Although UT is sometimes called "UTC" or "GMT" in other sources, this specification uses the term "UT" to avoid confusion with UTC or with GMT.

UNIX Time: The time as returned by the C `time()` function (see [Section 3](#) of the "System Interfaces" Volume of [\[POSIX\]](#)). This is an integer number of seconds since the POSIX Epoch (1970-01-01 00:00:00 UTC) not counting leap seconds. As an extension to POSIX, negative values represent times before the POSIX Epoch, using UT.

Wall Time: The time as shown on a clock set according to a location's law or practice.

3. The Time Zone Information Format (TZif)

The time zone information format begins with a fixed 44-octet header ([Section 3.1](#)) followed by a variable-length data block ([Section 3.2](#)) using four-octet (32-bit) transition times and leap second occurrences. These 32-bit values are limited to representing times from 1901-12-13 20:45:52 through 2038-01-19 03:14:07 UT.

The TZif header contains a field which specifies the version of the file's format. Version 1 files terminate after the 32-bit data block.

Version 2 and 3 files extend the format by appending a second 44-octet header, another variable-length data block using eight-octet (64-bit) transition times and leap second occurrences, and a variable length footer ([Section 3.3](#)). These 64-bit values can represent times approximately 292 billion years into the past or future.

A TZif file is structured as follows:

Version 1	Versions 2 & 3
+-----+	+-----+
Header for	Header for
32-bit	32-bit
Transitions	Transitions
+-----+	+-----+
Data with	Data with
32-bit	32-bit
Transitions	Transitions
+-----+	+-----+
	Header for
	64-bit
	Transitions
	+-----+
	Data with
	64-bit
	Transitions
	+-----+
	Footer
	+-----+

General Format of TZif Files

Version 1 files are considered a legacy format and SHOULD NOT be generated, as they do not support transition times after the year 2038.

Implementations SHOULD generate version 2 or 3 files. The sequence of time changes defined by the 32-bit header and data block SHOULD be a contiguous subsequence of the time changes defined by the 64-bit header and data block. When reading a version 2 or 3 file, implementations SHOULD ignore the 32-bit header and data block except for the purpose of skipping over them.

NOTE: All multi-octet integer values MUST be stored in network octet order format (high-order octet first, otherwise known as big-endian), with all bits significant. Signed integer values MUST be represented using two's complement.

[3.1.](#) TZif Header

The TZif header is structured as follows (the number of octets occupied by a field is shown in parenthesis):


```

+-----+-----+
| magic    (4) | ver |
+-----+-----+-----+
|           [unused - reserved for future use] (15)           |
+-----+-----+-----+-----+
| isutcnt  (4) | isstdcnt (4) | leapcnt  (4) |
+-----+-----+-----+-----+
| timecnt  (4) | typecnt  (4) | charcnt  (4) |
+-----+-----+-----+-----+

```

TZif Header

The fields of the header are defined as follows:

magic: The four-octet ASCII sequence "TZif" (0x54 0x5A 0x69 0x66) which identifies the file as utilizing the Time Zone Information Format.

ver(sion): An octet identifying the version of the file's format. The value **MUST** be one of the following:

NUL (0x00) Version 1 - The file contains only the 32-bit header and data block. Version 1 files **MUST NOT** contain a 64-bit header, data block, or footer.

'2' (0x32) Version 2 - The file **MUST** contain both the 32-bit header and data block and a 64-bit header, data block, and footer. The TZ string in the footer **MUST** strictly adhere to the expanded format of the "TZ" environment variable as defined in [Section 8](#) of the "Base Definitions" Volume of [\[POSIX\]](#).

'3' (0x33) Version 3 - The file **MUST** contain both the 32-bit header and data block and a 64-bit header, data block, and footer. The TZ string in the footer string **MAY** use extensions ([Section 3.3.1](#)) to the expanded format of the "TZ" environment variable as defined in [Section 8](#) of the "Base Definitions" Volume of [\[POSIX\]](#).

isutcnt: A four-octet unsigned integer specifying the number of UT/local indicators contained in the data block - **MUST** either be zero or equal to 'typecnt'.

isstdcnt: A four-octet unsigned integer specifying the number of standard/wall indicators contained in the data block - **MUST** either be zero or equal to 'typecnt'.

leapcnt: A four-octet unsigned integer specifying the number of leap second records contained in the data block.

timecnt: A four-octet unsigned integer specifying the number of transition times contained in the data block.

typecnt: A four-octet unsigned integer specifying the number of local time type records contained in the data block - MUST NOT be zero.

charcnt: A four-octet unsigned integer specifying the total number of octets used by the set of time zone designations contained in the data block.

3.2. TZif Data Block

The TZif data block consists of seven variable-length elements, each of which is series of zero or more items. The number of items in each series is determined by the corresponding count field in the header. The total length of each element is calculated by multiplying the number of items by the size of each item. Therefore, implementations that do not wish to parse or use the 32-bit data block can calculate its total length and skip directly to the header of the 64-bit data block.

In the initial data block, time values are 32-bit (`TIME_SIZE = 4` octets). In the second data block, present only in version 2 and 3 files, time values are 64-bit (`TIME_SIZE = 8` octets).

The data block is structured as follows (the number of octets occupied by a field is shown in parenthesis):

```

+-----+
| transition times          (timecnt x TIME_SIZE) |
+-----+
| transition types          (timecnt)              |
+-----+
| local time type records   (typecnt x 6)          |
+-----+
| time zone designations    (charcnt)              |
+-----+
| leap second records       (leapcnt x (TIME_SIZE + 4)) |
+-----+
| standard/wall indicators  (isstdcnt)             |
+-----+
| UT/local indicators       (isutcnt)              |
+-----+

```

TZif Data Block

The elements of the data block are defined as follows:

transition times: A series of four- or eight-octet UNIX time values sorted in strictly ascending order. Each value is used as a transition time at which the rules for computing local time may change. The number of time values is specified by the 'timecnt' field in the header.

transition types: A series of one-octet unsigned integers specifying the type of local time of the corresponding transition time. These values serve as indices into the array of local time type records. The number of type indices is specified by the 'timecnt' field in the header. Each type index MUST be in the range [0, 'typecnt').

local time type records: A series of six-octet records specifying a local time type. The number of records is specified by the 'typecnt' field in the header. Each record has the following format:

```
+-----+--+---+
|  utoff (4)   |dst|idx|
+-----+--+---+
```

utoff: A four-octet signed integer specifying the number of seconds to be added to UT in order to determine local time.

(is)dst: A one-octet value indicating whether local time should be considered Daylight Savings Time (DST). A value of one (1) indicates that DST is in effect. A value of zero (0) indicates that standard time is in effect.

(desig)idx: A one-octet unsigned integer specifying an index into the series of time zone designation characters, thereby selecting a particular designation string. Each index MUST be in the range [0, 'charcnt').

time zone designations: A series of ASCII characters constituting an array of NUL-terminated (0x00) time zone designation strings. The total number of characters is specified by the 'charcnt' field in the header. Note that two designations MAY overlap if one is a suffix of the other.

leap second records: A series of eight- or twelve-octet records specifying the corrections that need to be applied to UTC in order to determine TAI. The records are sorted by the occurrence time in strictly ascending order. The number of records is specified by the 'leapcnt' field in the header. Each record has one of the following structures:

32-bit Data Block:

```
+-----+-----+
| occur (4) | corr (4) |
+-----+-----+
```

64-bit Data Block:

```
+-----+-----+-----+
| occur (8) | corr (4) |
+-----+-----+-----+
```

occur(rence): A four- or eight-octet UNIX time value specifying the time at which a leap second correction occurs.

corr(ection): A four-octet signed integer specifying the total number of leap seconds to be applied to UTC on or after the occurrence. The correction values in adjacent leap second records **MUST** differ by exactly one (1).

standard/wall indicators: A series of one-octet values indicating whether the transition times associated with local time types were specified as standard time or wall clock time. A value of one (1) indicates standard time, and **MUST** be set to one (1) if the corresponding UT/local indicator is set to one (1). A value of zero (0) indicates wall time. The number of values is specified by the 'isstdcnt' field in the header. If 'isstdcnt' is zero (0), all transition times associated with local time types are assumed to be specified as wall time.

UT/local indicators: A series of one-octet values indicating whether the transition times associated with local time types were specified as UT or local time. A value of one (1) indicates UT, and the corresponding standard/wall indicator **MUST** also be set to one (1). A value of zero (0) indicates local time. The number of values is specified by the 'isutcnt' field in the header. If 'isutcnt' is zero (0), all transition times associated with local time types are assumed to be specified as local time.

The type corresponding to a transition time specifies local time for timestamps starting at the given transition time and continuing up to and not including the next transition time. Local time for timestamps before the first transition is specified by the first time type (time type 0). Local time for timestamps on or after the last transition is specified by the TZ string in the footer ([Section 3.3](#)) if present and nonempty, and is unspecified otherwise. If there are no transitions, local time for all timestamps is specified by the TZ

string in the footer if present and nonempty, and is specified by time type 0 otherwise.

A given pair of standard/wall and UT/local indicators is used to designate whether the corresponding transition time was specified as UT, standard time, or wall clock time. Note that there are only three combinations of the two indicators given that the standard/wall value **MUST** be one (1) if the UT/local value is one (1). This information can be useful if the transition times in a TZif file need to be transformed into transitions appropriate for another time zone (e.g. when calculating transition times for a simple POSIX TZ string such as "AKST9AKDT").

In order to eliminate unused space in a TZif file, every local time type record **SHOULD** be used by at least one transition (except that time type 0 need not be used if there is at least one transition). Likewise, every character in the time zone designations array **SHOULD** be used by at least one time type record.

3.3. TZif Footer

The TZif footer is structured as follows (the number of octets occupied by a field is shown in parenthesis):

```
+---+-----+---+
| NL|  TZ string (0...) |NL |
+---+-----+---+
```

TZif Footer

The elements of the footer are defined as follows:

NL: An ASCII new line character (0x0A).

TZ string: A rule for computing local time changes after the last transition time stored in the 64-bit data block. The string is either empty or uses the expanded format of the "TZ" environment variable as defined in [Section 8](#) of the "Base Definitions" Volume of [\[POSIX\]](#). If empty, the corresponding information is not available. If the string is nonempty, at least one transition **MUST** be present in the 64-bit data, and the TZ string **MUST** be consistent with the last 64-bit transition - i.e., evaluating the TZ string at the time of the last transition should yield the same time type as the time type specified in the last transition. Note that the string **MUST NOT** be NUL-terminated and **SHOULD NOT** begin with the ':' (colon) character.

3.3.1. TZ String Extensions

Version 3 TZif files MAY use the following extensions in the TZ string:

- o The hours part of the transition times may be signed and range from -167 through 167 instead of the POSIX-required unsigned values from 0 through 24.

Example: <-03>3<-02>,M3.5.0/-2,M10.5.0/-1

This represents a time zone that observes daylight saving time from 22:00 on the day before March's last Sunday until 23:00 on the day before October's last Sunday. Standard time is 3 hours west of UT and is abbreviated "-03"; daylight saving time is 2 hours west of UT and is abbreviated "-02".

- o DST is considered to be in effect all year if it starts January 1 at 00:00 and ends December 31 at 24:00 plus the difference between daylight saving and standard time, leaving no room for standard time in the calendar.

Example: EST5EDT,0/0,J365/25

This represents a time zone that observes daylight saving time all year. It is 4 hours west of UT and is abbreviated "EDT".

4. Use with the Time Zone Data Distribution Service

The Time Zone Data Distribution Service (TZDIST) [[RFC7808](#)] is a service that allows reliable, secure, and fast delivery of time zone data and leap second rules to client systems such as calendaring and scheduling applications or operating systems.

A TZDIST service MAY supply time zone data to clients in the Time Zone Information Format. Such a service MUST indicate that it supports this format by including the MIME media type "application/tzif" ([Section 7](#)) in its "capabilities" response (see [Section 5.1 of \[RFC7808\]](#)).

TZDIST clients MUST use the HTTP "Accept" [[RFC7231](#)] header field to indicate their preference to receive data in the "application/tzif" format.

4.1. Example

In this example, the client checks the server for the available formats and then requests that the time zone with a specific time zone identifier be returned in Time Zone Information Format.

Note that this example presumes that the time zone context path has been discovered (see [\[RFC7808\] Section 4.2.1](#)) to be "/tzdist".

>> Request <<

```
GET /tzdist/capabilities HTTP/1.1
Host: tz.example.com
```

>> Response <<

```
HTTP/1.1 200 OK
Date: Fri, 01 Jun 2018 14:52:23 GMT
Content-Type: application/json; charset="utf-8"
Content-Length: xxxx
```

```
{
  "version": 1,

  "info": {
    "primary-source": "IANA:2018e",
    "formats": [
      "text/calendar",
      "application/tzif",
    ],
    ...
  },
  ...
}
```

>> Request <<

```
GET /tzdist/zones/America%2FNew_York HTTP/1.1
Host: tz.example.com
Accept: application/tzif
```

>> Response <<

```
HTTP/1.1 200 OK
Date: Fri, 01 Jun 2018 14:52:24 GMT
Content-Type: application/tzif
Content-Length: xxxx
ETag: "123456789-000-111"
```

```
TZif2...[binary data]...
EST5EDT,M3.2.0,M11.1.0
```


5. Security Considerations

None.

6. Privacy Considerations

None.

7. IANA Considerations

This document defines a MIME [[RFC6838](#)] media type for the exchange of data utilizing the Time Zone Information Format.

Type name: application

Subtype name: tzif

Required parameters: N/A

Optional parameters: N/A

Encoding considerations: Binary

Security considerations: N/A

Interoperability considerations: N/A

Published specification: This specification.

Applications that use this media type: This media type is designed for widespread use by applications that need to exchange time zone information. The Time Zone Distribution Service [[RFC7808](#)] can directly use this media type.

Fragment identifier considerations: N/A

Additional information:

Magic number(s): The first 4 octets are 0x54, 0x5A, 0x69, 0x66

File extensions(s): N/A

Macintosh file type code(s): N/A

Person & email address to contact for further information:

Time Zone Database mailing list <tz@iana.org>

Intended usage: COMMON

Restrictions on usage: N/A

Author: See the "Author's Address" section of this document.

Change controller: IETF

8. Acknowledgments

The authors would like to thank the following individuals for contributing their ideas and support for writing this specification: Michael Douglass and Eliot Lear.

9. References

9.1. Normative References

[POSIX] IEEE, "Standard for Information Technology--Portable Operating System Interface (POSIX(R)) Base Specifications, Issue 7", IEEE 1003.1-2017, DOI 10.1109/IEEESTD.2018.8277153, January 2018, <<https://ieeexplore.ieee.org/servlet/opac?punumber=8277151>>.

This is identical to The Open Group Base Specifications Issue 7, 2018 edition [2].

[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>>.

[RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", [BCP 13](#), [RFC 6838](#), DOI 10.17487/RFC6838, January 2013, <<https://www.rfc-editor.org/info/rfc6838>>.

[RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", [RFC 7231](#), DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/info/rfc7231>>.

[RFC7808] Douglass, M. and C. Daboo, "Time Zone Data Distribution Service", [RFC 7808](#), DOI 10.17487/RFC7808, March 2016, <<https://www.rfc-editor.org/info/rfc7808>>.

- [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>>.

9.2. Informative References

- [RFC5545] Desruisseaux, B., Ed., "Internet Calendaring and Scheduling Core Object Specification (iCalendar)", [RFC 5545](#), DOI 10.17487/RFC5545, September 2009, <<https://www.rfc-editor.org/info/rfc5545>>.
- [RFC6557] Lear, E. and P. Eggert, "Procedures for Maintaining the Time Zone Database", [BCP 175](#), [RFC 6557](#), DOI 10.17487/RFC6557, February 2012, <<https://www.rfc-editor.org/info/rfc6557>>.
- [tz-link] Eggert, P. and A. Olson, "Sources for Time Zone and Daylight Saving Time Data", 2018, <<https://www.iana.org/time-zones/repository/tz-link.html>>.

9.3. URIs

- [1] <https://tools.ietf.org/html/bcp14>

Appendix A. Change History (To be removed by RFC Editor before publication)

Changes since -04:

- o Edited text discussing timestamps before first and after last transition.
- o Specified legal range of time type indices and time zone designation indices.
- o Notes that corrections in adjacent leap second records must differ by one.
- o Added recommendations to eliminate unused space.
- o Minor editorial changes.

Changes since -03:

- o Removed definition of GMT.
- o Updated definitions of UTC, TAI, and UT

- o Switched to using UT rather than UTC.
- o Added more text about the use of standard/wall and UT/local indicators.
- o Added Acknowledgments.
- o Minor editorial changes.

Changes since -02:

- o Updated definitions of Standard Time and DST.
- o Added definitions of GMT and UT.
- o Added a definition of Time Zone Data from [RFC7808](#).
- o Removed sentence stating that TZDB is accurate.
- o Added more text for standard/wall and UTC/local indicators and counts.
- o Added text discussing timestamps before first and after last transition.
- o Added more guidance text regarding 32-bit and 64-bit data consistency.
- o Minor editorial changes.

Changes since -01:

- o Renamed "POSIX Time" to "UNIX Time" and noted that values can be negative.
- o Noted that signed values MUST be represented using two's complement.
- o Renamed "POSIX TZ string" to "TZ string" and noted that it can be empty.
- o Moved TZ string extensions into its own subsection with examples.
- o Renamed leap second "epoch" to "occurrence".
- o Editorial changes from Paul Eggert.

Changes since -00:

- o Split TZif format description into a general overview and 3 subsections.
- o Updated Keywords boilerplate.
- o Updated POSIX reference.
- o Editorial changes from Eliot Lear.

Authors' Addresses

Arthur David Olson

Email: arthurdavidolson@gmail.com

Paul Eggert

University of California, Los Angeles

Email: eggert@cs.ucla.edu

Kenneth Murchison

FastMail US LLC

Email: murch@fastmailteam.com

