

Structs DSV Format

The XINA Structs DSV (delimiter separated values) format provides a standard delimited text data file format. This is recommended for data files attached to events, and forms the basis for the [structs buffer file format](#).

Files have certain standard requirements:

- Must be UTF-8 encoded
- New lines will be interpreted from either `\n` or `\r\n`
- Blank lines will be ignored
- Lines starting with the `#` character are treated as comments and ignored

The `conf` object may define other customization of the format:

| Key | Value | Default | Description |
|--------------|--|---|--|
| delimiter | string | auto detect (<code>'</code> , <code>\t</code> , <code>;</code>) | value delimiter |
| quote_char | character | <code>"</code> (double quote character) | value quote character |
| ignore_lines | number | <code>0</code> | lines to ignore at the start of the file |
| mode | <code>"row"</code> or <code>"col"</code> | auto-detect | file row/col format (see below) |
| t | <code>"auto"</code> , <code>"iso8601"</code> , <code>"s"</code> , <code>"ms"</code> , or <code>"us"</code> | <code>"auto"</code> | time format (see below) |
| zone | string | | time zone to use if not provided |
| invalid | <code>"ignore"</code> , <code>null</code> , or number | <code>"ignore"</code> | preferred interpretation of invalid literal |
| nan | <code>"ignore"</code> , <code>null</code> , or number | <code>"ignore"</code> | preferred interpretation of <code>'Nan'</code> literal |
| p_infinity | <code>"ignore"</code> , <code>null</code> , or number | <code>"ignore"</code> | preferred interpretation of positive <code>'Infinity'</code> literal |
| n_infinity | <code>"ignore"</code> , <code>null</code> , or number | <code>"ignore"</code> | preferred interpretation of negative <code>'Infinity'</code> literal |

It is strongly recommended to include a unique [appropriately generated 128-bit UUID in the standard 36 character format](#) as a comment in the first processed line of each file. (If `ignore_lines > 0`, this would be the first line after that number of lines.)

The first processed uncommented line will be interpreted as the column header. If the `mode` property is `"row"`, the file must contain three columns:

| Name | Description | Alternate Names |
|----------|---|----------------------------|
| t | Unix time or ISO8601 zoned timestamp | time, timestamp |
| k | key | key, mn, mnemonic, n, name |
| v | value (numeric, empty, or <code>null</code>) | val, value |

The header is used to determine the order of the columns.

For example (whitespace added for clarity, not required):

```
# 123e4567-e89b-12d3-a456-426614174000
t , k   , v
0 , v_mon , 1
0 , i_mon , 5
1 , t_mon , 100
2 , v_mon , 1.1
2 , i_mon , 4
3 , t_mon ,
4 , v_mon , 1.2
4 , i_mon , 3
5 , t_mon , 101
```

If `mode` is `"col"`, the file must first contain a time column, followed by a column for each mnemonic. The column headers must specify the mnemonic name or ID for each column. Unlike `row`, `null` values must be spelled out explicitly, as empty values will **not** create a point in the database.

For example, the following is equivalent to the above example (whitespace added for clarity, not required):

```
# 123e4567-e89b-12d3-a456-426614174000
t   , v_mon , i_mon , t_mon
0   , 1     , 5     ,
1   ,      ,      , 100
2   , 1.1   , 4     ,
3   ,      ,      , null
4   , 1.2   , 3     ,
5   ,      ,      , 101
```

If the `mode` property is not specified, the mode will be determined by the number of columns in the file. If there are exactly 3 columns with names matching the required columns for the `"row"` mode, that mode is used; otherwise the file is assumed to use the column mode.

Time Parsing

The mode of time processing is determined by the value for `t` in `conf`. The `auto` mode attempts to interpret the most likely formatting for the timestamp. If the value is an integer or floating point format, it will be interpreted as a Unix timestamp, with precision based on these rules:

- `t > 1e16`: error, value above typical range
- `t > 1e14`: microseconds
- `t > 1e11`: milliseconds
- `t > 1e8`: seconds

- `t <= 1e8`: error, value below typical range

Otherwise it will be interpreted as a zoned ISO8601 timestamp. If `t` is set explicitly in the configuration the time will always be interpreted in that context. The ISO timestamp may use the standard format: `2023-05-31T17:55:07.000` or condensed `20230531T175507.000`. If the `zone` property provided in the configuration, the timestamps do not require a zone. Otherwise they must include an explicit zone.

Revision #38

Created 28 July 2022 20:31:12 by Nick Dobson

Updated 16 September 2024 21:20:58 by Bradley Tse