

Exploring the Perl DateTime module

COUNTING OUT TIME

Because calendar rules are influenced by historical and political decisions, date manipulations are riddled with pitfalls. Perl's DateTime module knows all the tricks. **BY MICHAEL SCHILLI**

If a backup script launches at 10 pm and quits at 4 am, how long did it take to run? Six hours? Well, it depends. Just think about a process that ran between March 26 and 27 in 2005 somewhere in the UK. The clocks were put forward by one hour at 1 am, and that would make five hours the right an-

swer. If the same process had run at the same time in the USA, the answer would have been six hours, as summer time starts a week later in the US. But not in Indiana, which had not yet introduced summer time in 2005. In fact, Indiana is introducing summer time this year (2006, [2]). Fortunately, the DateTime

module ([5]) from CPAN knows all these historical and future rules and provides an easy interface to even the most complex date calculations.

What if you wanted to know how long the current summer time rules have been in use in the UK? Listing 1 (*dsthist*) discovers this by looking back from the year 2006 and checking through the month of March to find out if there is a day where you end up at 5 a.m. when adding three hours and one second to 00:59:59. If this happens, summer time was used during this year, and the script

Listing 1: dsthist

```
01 #!/usr/bin/perl -w
02 use strict;
03 use DateTime;
04
05 YEAR:
06 for my $year (
07 reverse 1964 .. 2006) {
08
09 for my $day (1 .. 31) {
10
11 my $dt = DateTime->new(
12 year => $year,
13 month => 3,
14 day => $day,
15 hour => 0,
16 minute => 59,
17 second => 59,
18 time_zone =>
19 "Europe/London",
20 );
21
22 $dt->add(
23 hours => 3,
24 seconds => 1
25 );
26
27 if ($dt->hour() == 5) {
28 print "$year: DST\n";
29 next YEAR;
30 }
31 }
32 print "$year: No DST\n";
33 last;
34 }
```

stops when it discovers that's not the case. The display shows that 1972 was the first year with today's summertime rules:

```
...
1974: DST
1973: DST
1972: DST
1971: No DST
```

Summer in the City

Europe has fairly uniform daylight saving time rules, but this is not true of the American continent. This not only applies to the various countries; even some US states do their own thing, and there are even a few counties that do not adopt the same approach as the states in which they are located. And to make things even more complicated, the rules have changed in the course of time.

Listing 2 (*dstchk*) uses *all_names()* to ascertain all the timezones known to the *DateTime::TimeZone* module (which is also available from CPAN). It jumps to the first of January in the timezone it is investigating and adds six months. If this returns a date with an hour value that is not equal to zero, some kind of time adjustment must have occurred in the first six months of the year, meaning that this timezone must have switched to summer time in this period.

Zones are normally stored in a "Continent/City" format; examples are *Europe/London* (for Great Britain), *Europe/Dublin* (for Ireland), *America/New_York* (the state of New York in the USA), *America/Vancouver* (the Canadian state of British Columbia), and *Pacific/Honolulu* (for Hawaii). But if a county has deviated from state rules at some time in the past, an additional subdivision becomes nec-

essary, for example, *America/Kentucky/Louisville* designates the US state of Kentucky, with its biggest city Louisville. (As you may be aware, Frankfort is the capital of Kentucky, although it is by no means the biggest town in the state.)

To reflect the fact that the county of Monticello in Kentucky was once (up to the year 2000) part of a timezone that is different from the one it is in now, *DateTime::TimeZone* has an entry for *America/Kentucky/Monticello*. Figure 1 shows the output from *dstchk*, showing daylight saving timezones in green using *Term::ANSIColor*.

DateTime can manipulate data for any timezone you like, including dates in the past, and timezones that have been through changes. Lord Howe Island just off the Australian coast has a quirky daylight savings time rule that puts the clock just half an hour forward or back. Listing 3 (*lord_howe*) shows that adding a second to 2005-10-30 01:59:59 returns a local time of 02:30:00.

A *DateTime* object created by the *new* constructor first exists in the special *floating* timezone, if the *time_zone* parameter doesn't specify the timezone explicitly. In this state, the timezone temporarily adapts to match other *DateTime* objects when calculations or comparisons are performed with them.

If you want to ignore daylight saving time in your time calculations, you can either use the "floating" zone, or you can choose the daylight-saving-time-free *UTC* (Universal Time Coordinated) timezone. Setting *time_zone* = > "local" for a *DateTime* object, to set the time to the zone in which your computer resides, forces *DateTime* to try out all kinds of tricks to guess your timezone configuration:

```
mschilli@mybox:~/DEV
America/Adak: yes
America/Anchorage: yes
America/Anguilla: no
America/Antigua: no
America/Araguaina: no
America/Argentina/Buenos_Aires: no
America/Argentina/Catamarca: no
America/Argentina/Cordoba: no
America/Argentina/Jujuy: no
America/Argentina/La_Rioja: no
America/Argentina/Mendoza: no
America/Argentina/Rio_Gallegos: no
America/Argentina/San_Juan: no
America/Argentina/Tucuman: no
America/Argentina/Ushuaia: no
America/Aruba: no
America/Asuncion: yes
America/Bahia: no
America/Barbados: no
America/Belem: no
America/Belize: no
America/Boa_Vista: no
America/Bogota: no
America/Boise: yes
America/Cambridge_Bay: yes
America/Campo_Grande: yes
America/Cancun: yes
America/Caracas: no
America/Cayenne: no
America/Cayman: no
America/Chicago: yes
America/Chihuahua: yes
America/Coral_Harbour: no
America/Costa_Rica: no
America/Cuiaba: yes
America/Curacao: no
America/Danmarkshavn: no
America/Dawson: yes
America/Dawson_Creek: no
America/Denver: yes
America/Detroit: yes
America/Dominica: no
America/Edmonton: yes
America/Eirunpe: no
America/El_Salvador: no
America/Fortaleza: no
America/Glace_Bay: yes
America/Godthab: yes
America/Goose_Bay: yes
America/Grand_Turk: yes
America/Grenada: no
America/Guadeloupe: no
America/Guatemala: no
America/Guayaquil: no
America/Guyana: no
America/Halifax: yes
America/Havana: no
mybox:mschilli@eg*
```

Figure 1: Have you ever wanted a quick report on which areas of the American continent use daylight saving time?

Listing 2: dstchk

```
01 #!/usr/bin/perl -w
02 use strict;
03 use DateTime;
04 use Term::ANSIColor;
05 qw(:constants);
06
07 for my $zone (DateTime::
08     TimeZone::all_names()) {
09     my $from =
10         DateTime->now(
11             time_zone => $zone);
12
13     $from->truncate(
14         to => "year");
15     my $to =
16         $from->clone()
17         ->add(months => 6);
18
19     print "$zone: ";
20
21     if ($to->hour() == 0) {
22         print RED, "no", RESET,
23             "\n";
24     } else {
25         print GREEN, "yes", RESET,
26             "\n";
27     }
28 }
```

Listing 3: lord_howe

```

01 #!/usr/bin/perl -w          13 'Australia/Lord_Howe',
02 use strict;                14 );
03 use DateTime;              15
04                               16 $dt->add(
05 my $dt = DateTime->new(     17 DateTime::Duration->new(
06   year    => 2005,          18   seconds => 1
07   month   => 10,           19 )
08   day     => 30,           20 );
09   hour    => 1,            21
10   minute  => 59,           22 # 2005-10-30 02:30:00
11   second  => 59,           23 print $dt->date(), " ",
12   time_zone =>              24 $dt->hms(), "\n";

```

```

my $dt = DateTime->now(
    time_zone => "local");
print $dt->time_zone()->name();

```

This returned *America/Los_Angeles* on a machine located in our Perlmeister lab in San Francisco. Not bad.

Rotational Drag

How many seconds elapsed between 00:59:00 and 01:00:00 January 1 1999? One minute, that is 60 seconds? Wrong! Tip: At this point in time, the standard timezone, UTC, which is identical with the Greenwich timezone, added an extra *leap second*, which lead to a 61-second minute!

As the site at [3] tells you, a second is no longer defined as a fraction of one day, and hasn't been since 1967, but as a

function of the far more constant resonance of the caesium 133 atom. The rotational speed of the earth has slowly decreased over the last 40 years. As the earth now takes slightly longer than 24 times 3600 atomic seconds for a single rotation, a leap second has been added here and there to official timekeeping since 1972, every 1.5 years on average. June 30 and December 31 are the reference dates. If official time is off by about one second, the second is added at the end of these days. This said, the earth seems to be back up to speed; 2005 was

the first year since 1998 to need a one-second boost. If the earth were to rotate more quickly, the strict timekeepers would simply subtract a leap second from official time. However, this has never happened in the brief history of leap seconds.

Listing 4 (*leapsec*) starts in the year 1960 and moves gradually up to the year 2005, searching for leap seconds on June 30 or December 31. To do so, it sets the time in the UTC timeone to 23:59:00, adds 60 seconds, and checks if an unusual value of 60 is returned for the seconds segment. If so, the minute in question must have had 61 seconds, and we have discovered a leap second. After completing this search, the timezone is set to "Europe/London" by calling *set_time_zone()*. A call to *print* outputs the local time and the number of leap seconds found so far.

Figure 2 shows the output. The different local leap times on July 1 are a result of historical daylight saving time changes.

Pot-Holed Abstractions

Unfortunately, legacy timekeeping on Unix systems, which counts the seconds since 1970, does not take leap seconds into consideration. Whereas the second

Listing 4: leapsec

```

01 #!/usr/bin/perl -w          18           minute => 59,
02 #####                      19           second  =>  0,
03 # leapsec - Print years with 20           time_zone => "UTC";
04                               21
05 # Mike Schilli, 2005         22   my $later =
06   (m@perlmeister.com)       23   $now->clone()->add(
07                               24       seconds => 60);
08                               25
09 my $secs;                   26   $later->set_time_
10                               27   zone("Europe/London");
11                               28
12 for my $year (1960..2005) {  29   if($later->second() == 60)
13   for my $date ([30,6],      30   {
14     [31,12]) {              31     print $later->dmy(), "
15     my $now = DateTime->new(  32     ",
16       year    => $year,      33     $later->hms(), ":
17       month   => $date->[1],  34     ++$secs, "\n";
18       day     => $date->[0],  35   }
19       hour    => 23,         36   }

```

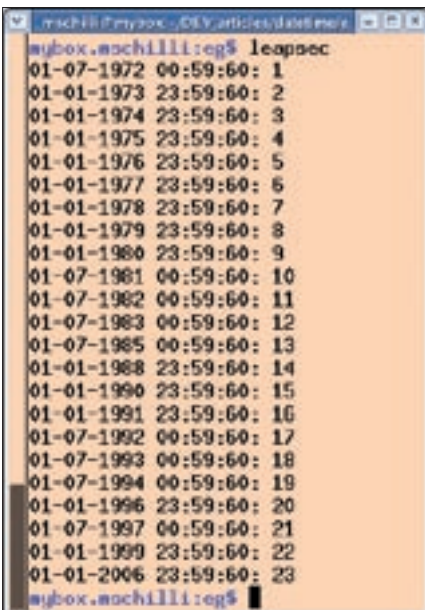


Figure 2: Leap seconds from 1960 to today.

Listing 5: leapreveal

```

01 #!/usr/bin/perl -w
02 use strict;
03 use Sysadm::Install qw(:all);
04
05 use DateTime;
06
07 my $dt = DateTime->new(
08     year      => 1990,
09     time_zone => 'UTC'
10 );
11
12 $dt->add(
13     seconds => 3600 * 24 *
14     5000);
15 print "$dt\n";

```

hand moved from 23:59:59 to 23:59:60 December 31, 1998, in the UK, the counter on Unix machines that follow the POSIX standard moved from 915148799 to 915148800. The next virtual hop from 00:59:60 to 01:00:00, however, wasn't reflected by any Unix time counter, both points of time are correctly represented by a Unix time value of 915148800.

If you subtract two Unix times from one another, and calculate the UTC time that has elapsed between them, you may need to correct the results if a leap second has occurred between the two dates. For more details of this confusing approach, see [6] and [7].

DateTime provides the class method *from_epoch*(*epoch* => *\$time*), which constructs a *DateTime* object from a Unix counter. The *epoch*() method of a *DateTime* object does the opposite, returning a counter value.

Listing 5 (*leapreveal*) shows what happens if you simply add the number of seconds in 5000 days to the date 1.1.1990: the result of this calculation is 2003-09-09T23:59:53. In other words, the answer reveals that there are 7 seconds missing from the end of the day, caused by leap seconds in between the two dates! On the other hand, using *add*(*days* => 5000) to add 5,000 days returns a result of 2003-09-10T00:00:00. *DateTime* strictly separates the handling of time units such as days and seconds and will not normally convert a time value such as "5000 days" to seconds.



Figure 3: Different languages and customs applied to converting and formatting a date string.

But if you want to convert days to seconds, and will excuse the following potted abstraction, you can use the *\$to->subtract_datetime_absolute(\$from)* method to subtract the *DateTime* object *\$from* from a *DateTime* object *\$to* to obtain a *DateTime::Duration* object, and this object's *seconds()* method really does give the exact number of seconds that elapsed during the period.

Superman

The dateline is another curiosity. If you fly east, the local time in the timezones you fly through gets later and later. At some point, the date has to shift to the previous day. If this were not so, you would be able to travel to the future in a fast plane. The dateline ([4]) crosses the Pacific from north to south, slightly to the east of the island groups off the South East Asia shore.

Listing 6 (*daytrip*) shows what happens if you take a fast plane from Japan

(west of the date line) to Hawaii (east of the date-line). These are your flight details:

```

Departure: Sunday, 2
January 29 2006, 07:30
Arrival: Saturday, 2
January 28 2006, 19:00

```

In other words, you leave on Sunday morning and arrive a day earlier, although the flight takes six-and-a-half hours. Saturday evening before the lottery results are announced – pity this only works for local time.

Speaking in Tongues

Listing 6 (*daytrip*) also shows how *DateTime* handles different date formats. It uses formatting tools from the *DateTime::Format::** class hierarchy both to parse a date string with *parse_datetime()*, and to output the results. *DateTime::Format::Strptime* is a particularly flexible formatting tool that has placeholders for the format string, following a similar approach to the *strptime()* function in C. *%A* represents the weekday, *%B* the month name in writing, *%d* the date, *%H* the hour, and so on. The *locale* parameter is set to "en_GB" for Great Britain. *en* selects the English language. Part two of the locale specifies the country and its special rules.

Listing 6: daytrip

```

01 #!/usr/bin/perl -w                                07:30"
02 use strict;                                       19 );
03 use DateTime;                                     20
04 use                                               21 $dt->set_formatter($format);
05     DateTime::Format::Strptime;                 22
06                                                     23 print "Departure: $dt\n";
07 my $format =                                       24
08     DateTime::Format::Strptime                 25 $dt->add(
09     ->new(                                       26     DateTime::Duration->new(
10     pattern =>                                   27     hours => 6,
11     "%A, %B %d %Y, %H:%M",                       28     minutes => 30
12     locale => "en_UK",                            29 )
13     time_zone => 'Asia/Tokyo',                    30 );
14 );                                               31
15                                                     32 $dt->set_time_zone(
16 my $dt =                                         33     'Pacific/Honolulu');
17     $format->parse_datetime(                       34 print "Arrival: $dt\n";
18     "Sunday, January 29 2006,

```

Listing 7: locales

```

01 #!/usr/bin/perl -w          14
02 use strict;                15 $dt->set_locale($locale);
03 use DateTime;              16
04 use                        17 my $format =
05   DateTime::Format::Strptime; 18   DateTime::Format::Strptime
06                               19   ->new(
07 my $dt = DateTime->now();    20   pattern => $dt->locale()
08                               21   ->long_datetime_format()
09 for my $locale (           22   );
10 qw(en_AU en_US de_DE fr_FR 23
11 es_ES es_MX)              24 $dt->set_formatter($format);
12 )                          25 print "$locale: $dt\n";
13 {                            26 }

```

In Listing 7 (*locales*) are more examples: *en_AU* and *en_US* are locales for Australian and US English. *fr_FR* selects French; *es_ES*, and *es_MX* give you Spanish for Spain and Mexico. After initializing the formatter with an appropriate locale value, it is passed to the *DateTime* object using *set_formatter*. “Stringified” *DateTime* objects are converted to strings. Figure 3 shows a few examples.

Leap Years

Looking back to the year 2000, it is probably safe to assume that most program-

mers are aware of the rule that a leap year occurs every four years, but not if the year is divisible by 100, the exception being years that are divisible by 400.

Of course, *DateTime* understands these rules, so just to prove a point, let’s tackle a more complex problem: How long is the list of Friday the 29ths of February between 1980 and 2000?

Two sets of *DateTime* objects give us an elegant approach to solving this problem of finding the Friday leap days: we need to store all the Fridays in one of them, and all the 29ths of February in

Listing 8: frifeb29

```

01 #!/usr/bin/perl -w          19
02 use strict;                20 my $it = $set->iterator(
03 use DateTime;              21   start => DateTime->new(
04 use                        22   year => 1980
05   DateTime::Event::Recurrence; 23   ),
06                               24   end => DateTime->new(
07 my $feb29 =                25   year => 2020
08   DateTime::Event::Recurrence 26   ),
09   ->yearly(                 27   );
10   days => 29,               28
11   months => 2,              29 while (my $dt = $it->next())
12   );                        30 {
13 my $fri =                  31   $dt->set_locale("en_GB");
14   DateTime::Event::Recurrence 32   print $dt->day_name(), " ",
15   ->weekly(days => 5);      33   $dt->month_name(), " ",
16                               34   $dt->day(), " ",
17 my $set =                  35   $dt->year(), "\n";
18   $fri->intersection($feb29); 36 }

```

the other. A *DateTime::Set* class object can theoretically contain an infinite number of *DateTime* objects.

The easiest way to create a set for this purpose is to use the CPAN *DateTime::Event::Recurrence* module. The *DateTime::Event::Recurrence->yearly(days => 29, months => 2)*; constructor gives us a *DateTime::Set* type object with all the 29ths of February as an abstract description.

At the same time, Listing 8 (*frifeb29*) uses *weekly(days => 5)* to define a second set that contains all Fridays (that is, all the 5th days in the week). We can then use the *intersection()* method to give us a set of Friday the 29ths.

To tell the iterator (defined in line 20) for the resulting set where to start, the *start* parameter defines the starting date, and a *DateTime* object set to the year 2020 is used to set the end date.

The *while* loop in line 29 uses *next()* for kicking off the iterator and pushing it towards the end of the period of time under review. The result of this investigation is: there was a Friday, February 29th in 1980, and there will be another one in 2008. ■

INFO

- [1] Listings for this article:
<http://www.linux-magazine.com/Magazine/Downloads/64/Perl>
- [2] “What time is it in Indiana?“,
<http://www.mccsc.edu/time.html>
- [3] Leap seconds:
http://en.wikipedia.org/wiki/Leap_second
- [4] The date line: http://en.wikipedia.org/wiki/International_Date_Line
- [5] Datetime project homepage:
<http://datetime.perl.org>
- [6] Unix Time:
http://en.wikipedia.org/wiki/Unix_time
- [7] UTC, TAI, and UNIX time:
<http://cr.yip.to/proto/utctai.html>

THE AUTHOR

Michael Schilli works as a Software Developer at Yahoo!, Sunnyvale, California. He wrote “Perl Power” for Addison-Wesley and can be contacted at mschilli@perlmeister.com. His homepage is at <http://perlmeister.com>.

