

Geoff Huston  
December 2016

## Leaving it to the Last Second

Thanks to the moon, the earth's rate of rotation is slowing down. It's a subtle interaction and the modelling of planetary dynamics predicts that the earth's rotation should be slowing down by an average of 2.3 milliseconds per century. But this is not quite so uniform, as the Economist reported in December on work of analysing ancient records of solar eclipses:

"After crunching the numbers, the team found that the actual rate at which days have been lengthening over the past couple of millennia is 1.8 milliseconds per century, considerably slower than the 2.3 milliseconds predicted. The main reason for the difference, they reckon, is the lingering effect of the most recent ice age, during which the mass of ice at the planet's poles was sufficient to deform its shape and thus alter its rate of spin. That is not, though, the only thing which is happening. The researchers also found small but cyclical patterns in the rate of change that repeat themselves over decades—as well as intriguing hints of longer cycles with time periods of thousands of years. Exactly what geophysical goings-on such cycles represent is one for the geologists to work out."

<http://www.economist.com/news/science-and-technology/21711302-value-ancient-records-ancient-eclipses-show-how-days-are-getting-longer>

So what?

Maybe we can start by looking at how we've defined time over history.

The ancient Egyptians divided a day into 12 night hours and 12 daylight hours, adjusting the duration of each of these hours according to the season, so that nights and days were always divided into exactly 12 hours. The ancient Babylonians used a sexagesimal system, dividing time into successively finer units using a base of 60. It appears that it was the Arabs at the time of the first millennium who melded the two systems together, giving us a day composed of 24 hours, but dividing each hour using the Babylonian system into 60 minutes and each minute into 60 seconds. Each hour, minute and second was of constant duration. The imperceptible slowing of the earth's rotation meant nothing at the time, so that we were happy to stretch out what we accepted as a second of time such that it was always 1/86,400 of a single rotation of the earth about its own axis (the sun is intended to be at the same position relative to the earth after each rotation, which is why this time base is termed "solar time").

It has only been in recent decades that we've turned our attention to timekeeping with an obsessive level of detail that rivals, and maybe even surpasses, train spotting. And this simple sun-based time base has been unable to withstand the pressure of such intense horological scrutiny. In 1955 we went down the route of a standard definition of a "second" as a fixed time interval, so that a "second" always spanned the same time duration.

This statement about the constancy of a time standard just begs the noting of the changes to this assumption introduced by the theory of relativity, where both space and time are distorted in order to preserve

the constancy of the speed of light, irrespective of observers' reference frames. But for our purposes let's stick to a Newtonian world and assume a single inertial frame of reference where time beats at a constant invariant rate.

This "standard second" was defined by the International Astronomical Union as  $1/31,556,925.9747$  of the 1900 "mean tropical year". This definition was also adopted in 1956 by the International Committee for Weights and Measures and in 1960 by the General Conference on Weights and Measures, becoming a part of the International System of Units (SI). This definition addressed the problem of the drift in the value of the mean solar year, as well as the slowing of the earth's angular momentum by specifying a particular year as the baseline for the definition. However, by the mid 1960's even this definition was found to be inadequate for precise time measurements. In 1967 the SI second was again redefined, this time in experimental terms as a repeatable measurement. The new definition of a second was 9,192,631,770 periods of the radiation emitted by a caesium-133 atom in the transition between the two hyperfine levels of its ground state. This time base is "atomic time" (TAI).

All this is fine, except that the subtle influence of the earth's rotational slowing means that over the years our stunningly accurate clocks will start to drift away from the solar time baseline. The folk who found this to be an unacceptable deviation needed to find a way to accommodate both a fixed length second and the slowing rotation of the earth. The chosen solution was to keep the atomic time base (TAI) in sync with the earth's rotational solar time (Universal Coordinated Time, or UTC) by adjusting UTC time in units of 1 second from time to time. This is normally done when UTC drifts away from TAI by more than 0.6 second, and the "correction" is to insert, or remove, one entire SI-standard leap second into UTC time.

The UTC standard allows such correctional leap seconds to be applied at the end of any month, but since 1972 all of these leap seconds have been inserted either at the end of June 30 or December 31, making the final minute of the month in UTC either one second longer or one second shorter when a leap second is applied. The Earth Orientation Centre of the International Earth Rotation and reference System Service (IERS) (<http://hpiers.obspm.fr/eop-pc/index.php?index=webservice>) publishes announcements every six months, advising whether leap seconds are to occur or not in the forthcoming six months. Longer periods of advance notice are not possible because of the degree of uncertainty in predicting the precise value of the cumulative effect of fluctuations of the deviation of the Earth's rotational period from the value of the mean solar day.

Between 1972 and 2012 some 25 leap seconds have been added to UTC. On average this implies that a leap second has been inserted about every 19 months. However, the spacing of these leap seconds is quite irregular: there were no leap seconds in the seven-year interval between January 1, 1999 and December 31, 2005, but there were 9 leap seconds in the 8 years 1972–1979. Since December 31 1998 there have been only 4 leap seconds, on December 31 2005, December 31 2008, June 30 2012 and June 30, 2015, each of which have added one second to that final minute of the month, at the UTC time of day.

Leap second-wise, it's now crunch time again.

On Saturday 31st December 2016, the last minute of 2016 in UTC time will be extended to be 61 seconds long.

Now all this would be just arcane esoterica for most of us were it not for the fact that we've trained our computers to be horologically obsessive, so now pretty much every connected computer runs some form of time alignment protocol, such as the Network Time Protocol (NTP), to keep its clock in relatively good alignment with UTC.

NTP does not replace a system's internal time of day clock. Instead, it keeps the system's clock aligned to what NTP believe is the precise time of day by making periodic small adjustments to the system's clock. In this model, a leap second displaces the system's clock by one second from the NTP time, and NTP will subsequently shift the local clock back into sync with NTP's time by a sequence of clock adjustments. Alternatively, NTP can signal the leap second to the system clock directly and the system clock can make an immediate adjustment by winding the system's clock backwards by one second (as is done in POSIX systems). Google use a slightly different method, and use a 20 hour period before and after midnight UTC to stretch the direction of the signalled second to be 13.9 $\mu$ s slower than the SI second. This "smears" the leap second across 20 hours, and keeps individual changes within the realm of the existing clock corrections applied to the internal quartz oscillators (<https://developers.google.com/time/smear>). Other approaches to this leap second are documented by Red Hat (<https://access.redhat.com/articles/15145>) and Microsoft (<https://support.microsoft.com/en-us/kb/2722715>). An excellent analysis of what happened last time can be found in an article by Steven Strowes, Preparing for the 2016 Leap Second ([https://labs.ripe.net/Members/stephen\\_strowes/preparing-for-the-2016-leap-second](https://labs.ripe.net/Members/stephen_strowes/preparing-for-the-2016-leap-second)).

The POSIX approach of winding the kernel clock backward by one second to replay the last second of the minute appears to be the cause of many system crashes. Time-related code rarely assumes that the system time runs backwards, and can get horribly confused when it does. Smearing the change across a longer time interval makes some sense but different systems may perform this operation at different speeds. The objective of NTP is to create an environment where a collection of systems are synchronised in time to a relatively high degree of accuracy. During the period of the leap second smearing this assumption may not hold, and the collection of systems may differ in their view of the current UTC time by up to one second. Code might react badly the accuracy of a common time base drifts from milliseconds to an entire second.

Previous leap seconds have caused some quite impressive system crashes as a result of unfounded assumptions about the constancy of UTC time. And it seems that the problems are persistent, in that the leap second in June 2015 managed to generate a few visible crashes despite widely reported negative experience from a similar leap second event just three years earlier. From mid-2015 there were press reports of interruptions of around 40 minutes duration that occurred with Twitter, Instagram, Pinterest, Netflix, Amazon, and Apple's music streaming series Beats 1.

It would be nice to think that recent leap seconds have been a learning experience and that this time around the leap second at the end of 2016 will take place without a murmur from our computers. It would be nice to think that operating systems, applications and services will all quietly cope with this additional beat added to the regular pulsing of the heartbeat of time, whether its performed by slewing the local clock or repeating that last second. Nice, but overly naive. It's more likely to expect some services to once again experience a hiccup as the leap second occurs. Doubtless not all folk will be paying attention over the New Year weekend and some outages will occur.

If you operate services in the Internet you may want to check for yourself when midnight UTC rolls around that your online services are not quietly choking over this extra time bonus that has been stuffed literally into the last minute of 2016.

---

## Author

*Geoff Huston* B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for building the Internet within the Australian academic and research sector in the early 1990's. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001 and chaired a number of IETF Working Groups. He has worked as an Internet researcher, as an ISP systems architect and a network operator at various times.

*[www.potaroo.net](http://www.potaroo.net)*

---

## Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.