PLANES WILL CRASH! THINGS THAT LEAP SECONDS DIDN'T, AND DID, CAUSE

Steven L. Allen *

In 1970, just as the CCIR decreed that radio broadcast time signals would have leap seconds, IAU Commission 31 (Time) made its report to the 14th General Assembly in Brighton. Along with other objections the section on "Coordinated Time" indicated that "the world-wide collision avoidance system for aircraft (CAS) ... cannot admit stepping time adjustments". Four decades and 25 leap seconds later no planes have crashed, but computer operating systems have crashed. This presentation gives some looks into the recent news about the effects of leap seconds.

"REGRETTABLE MISUNDERSTANDINGS"

Finding the origins of the leap second is not an easy exercise. Most of the discussions occurred in the study groups and working parties of the International Radio Consultative Committee (CCIR), and the few records from those meetings are are hard to find. One of the best publicly available views of the process was included as part of a report to the General Assembly of the International Astronomical Union (IAU) in 1970.

The report from IAU Commission 31 (Time)[1] appears to have been assembled prior to the CCIR Plenary Assembly which approved Recommendation 460[2] and the inception of leap seconds. In the report is a section "Coordinated Time" which summarized the recent CCIR meetings on the subject of radio broadcast time signals. Sentences attributed in part to D.H. Sadler (who later explored his great dissatisfaction with radio broadcasts based on atomic time[3]) indicated that "Commissions ignore the needs of a great minority". They also gave the dire prediction that "the world-wide collision avoidance system for aircraft (CAS) ... cannot admit stepping time adjustments," i.e., that the institution of leap seconds might cause planes to crash.

The proceedings of the 1970 IAU General Assembly include the minutes from multiple agitated meetings of Commission 31[4] and a special Joint Meeting on Time Scales.[5] Some of the IAU members had been at the earlier CCIR Plenary Assembly which had approved Recommendation 460 with its requirement that leap seconds would begin as of 1972. Aside from pointing out that the decision to implement leap seconds had already been made, they also noted that the IAU had no basis to make an official response because the CCIR had failed to follow their own requirement to send a letter informing the IAU of the decision. Further insight into the unhappiness surrounding the institution of about the leap second came in 1972 when J. Terrien, the director of the International Bureau of Weights and Measures (BIPM), wrote "in the last few years regrettable misunderstandings, especially between astronomers and physicists, have crept into discussions on time and frequency."[6]

* UCO/Lick Observatory, 1156 High Street, Santa Cruz, CA 95064

In 1978 the report from CCIR Study Group 7 to the CCIR plenary assembly[7] drew attention to recommendations from multiple international agencies that UTC, with its leap seconds, had properties that allowed it to be adopted as the basis for civil time. The SG7 report also pointed out that some countries had already redefined their legal time to be based on UTC and wrote "the UTC time scale is the general reference for civil time". Thus SG7 recognized that the use of UTC had already accumulated political implications without clearly realizing that those would eventually compete with technical considerations.

The 1978 CCIR plenary assembly approved Recommendation 535 which indicated that the international treaty text of the Radio Regulations should adopt UTC as its basis for expressing time and date. Also at that meeting Study Group 7 abolished its interim working parties, including the one which had spent a decade creating and revising the scheme of leap seconds.[8]

The efforts of CCIR Study Group 7 to define UTC with leap seconds continued to spread into other official international documents. In 1979 the World Administrative Radio Conference of the ITU adopted a new version of the Radio Regulations.[9] The ITU wrote "UTC is equivalent to mean solar time at the prime meridian (0° longitude), formerly expressed in GMT." The term UTC is used throughout the 984 pages of the Final Acts, replacing GMT in earlier versions while still retaining an explicit statement that UTC is related to longitude and earth rotation. In 1980 the plenary assembly of the CCITT accepted UTC as "the time scale for all other telecommunication activities."[10]

None of the agencies that adopted UTC demonstrated foresight that the 1980s would be the decade when the operators of Unix computing systems began to realize the need to standardize their interfaces. The CCIR considered that it had completed all work for defining an international time scale, but at that time the text of CCIR recommendations was not freely available. The effect of CCIR Recommendation 460 was to define the only instance of a time scale where the duration of one second is unrelated to the duration of one day. The implications of this aspect of UTC were not evident to the committees who evolved standards for computing systems.

When prescribing the counting of elapsed seconds the computing committees decided that simplicity in the algorithms was more important than setting up a robust scheme for correct handling of the UTC time scale. The resulting standards for computing systems[11,12] were not consistent with the implementation of leap seconds in radio broadcast time signals. The information processing standards have remained inconsistent with TF.460 during their evolution over subsequent decades.[13,14,15,16,17] The result is that POSIX requires an operating system not to allow any of the running applications to be aware of a leap second, and different systems attempt this in different ways. To this date the practical effect of the POSIX standard is to require that the default system clock in Unix must count the mean solar seconds of earth rotation rather than the SI seconds of TAI.

2005: A LEAP SECOND AFTER SEVEN YEARS WITHOUT

During the 1990s and into the 2000s the rotation of the earth accelerated notably. The IERS announced fewer insertions of leap seconds because mean solar time was not falling as quickly behind the atomic chronometers. There were no leap seconds from 1998 July through 2005 June. Among other changes the same interval saw tremendous evolution of computing hardware, internet connectivity, and the end of Selective Availability degrading the signals from GPS satellites.

The maritime regulations for many large harbors have required that all large vessels employ an Automated Identification System (AIS) unit at any time that they are not anchored. AIS units consist of a GPS receiver which determines the position of the ship and a VHF transmitter which broadcasts

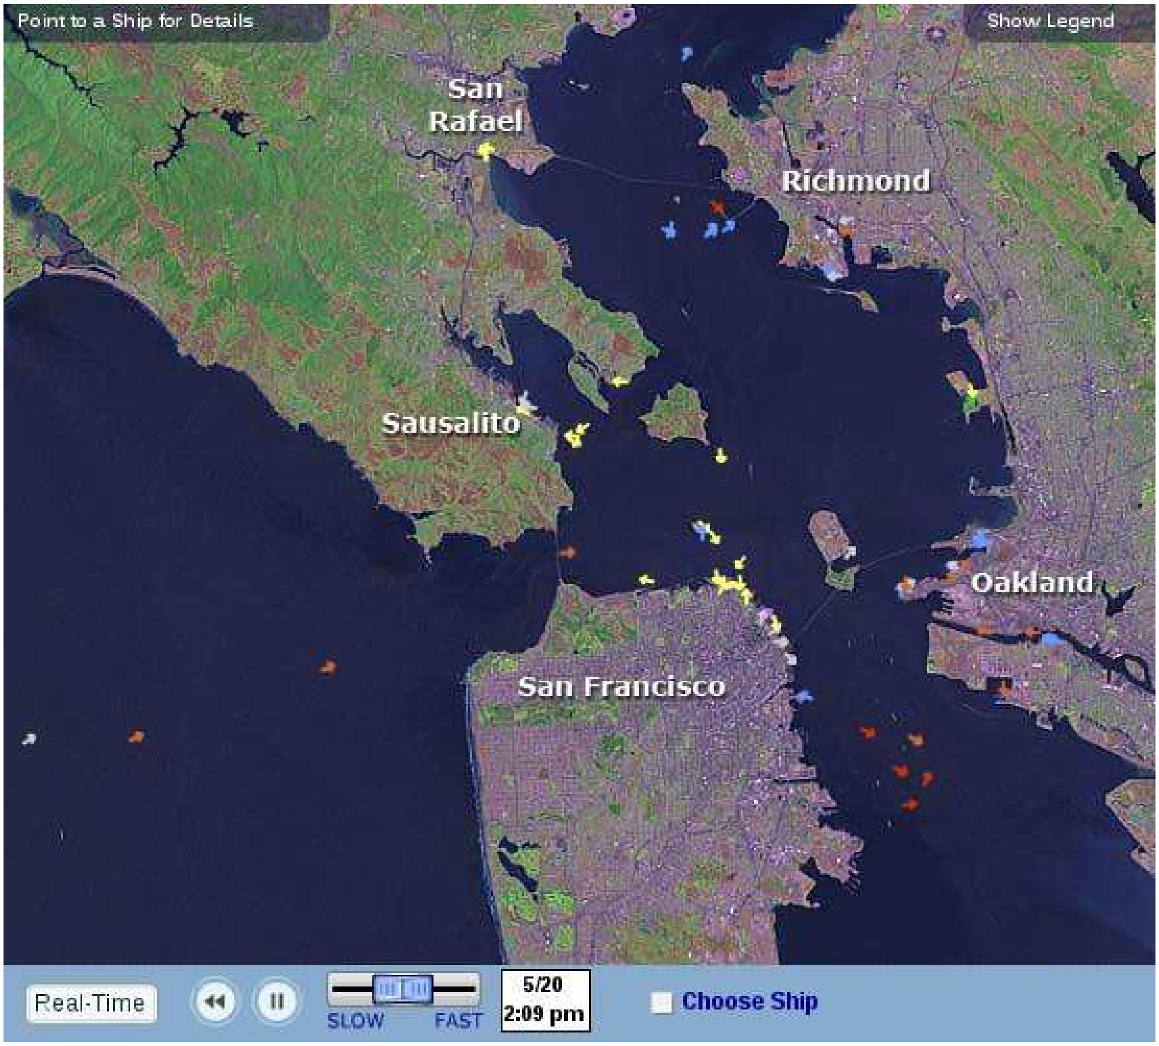

Figure 1. AIS map of San Francisco Bay from BoatingSF.com

that along with an identification of the vessel. AIS enables collision avoidance for the ships as well as tracking by anyone around the harbor. Figure 1 gives an AIS view of San Francisco Bay.* The last half of 2005 saw press releases from two different manufacturers of AIS units.

In September Saab informed customers about two models of its AIS systems.† According to their press release the software in those AIS units applied a leap second immediately upon receiving the GPS signals which announced the leap which should be inserted at the end of the year. They admitted the incorrect time in their AIS might cause signal collisions and that some other AIS units might not receive them. They indicated that no users had reported a problem, and they had a software upgrade which would fix the units without waiting for the new year.

*`http://www.boatingsf.com/ships-on-SF-bay/central-san-francisco-bay`

†Saab TransponderTech AB (2005-09-09), Timing problem in the R3 and R4 Transponder, Document Id PT-05-0078, issued by Mikael Pettersson and approved by Gunnar Mangs

In December ACR Electronics issued a press release* and technical bulletin† about the GlobalWatch and GlobalWatch2 AIS units. The text of both documents contained hints of blame at the earth and the authorities responsible for implementing TF.460: "For the first time since 1998, the International Earth Rotation and Reference Systems Service in Paris will sneak in an extra second at the end of the year". The technical bulletin indicated that after the leap second those AIS units might stop seeing other ships on their displays. It gave detailed instructions about two different manual procedures that could be used to restore proper function.

2008: A LINUX KERNEL DEADLOCK BUG

At the new year of 2009 Slashdot reported that Linux admins found crashed systems.‡ The telltale came from other Linux machines which did not crash; they had made a log entry `Clock: inserting leap second 23:59:60 UTC`. In a post to the Linux Kernel Mailing List on January 1 John Stultz confirmed that the bug occurred in some versions of the kernel where calling the `printk()` routine to make that log entry caused the deadlock to happen.§

2011: GOOGLE'S SOLUTION TO THE POSIX DILEMMA

In 2011 September site reliability engineer Christopher Pascoe posted a blog entry (Figure 2) about the effects of the previous two leap seconds on the internet cloud systems of Google.¶ This revealed that Google had seen problems with some of the machines in their cloud during the 2005 leap second. Google engineers looked at the amount of code that might be affected by a leap second, and they decided to adopt a strategy other than fixing all the code.

The solution that Google chose was to rely solely on their own internal network of Network Time Protocol (NTP) servers. Instead of inserting a leap second all at once they adjusted the frequency of their NTP servers using what is now called the Google "leap smear". At the end of 2008 the Google servers did not see the leap second and services offered by their cloud kept working.

Google admits that this solution is a "lie" which does not conform to the NTP specification, but implementing it saved them "massive amounts of time and energy". "The team involved in solving this issue was a handful of people, distributed around the world, who were able to work together without restriction in order to solve this problem." Google is one of few agencies capable of setting up, testing, and maintaining its own global time distribution scheme. Furthermore, the services offered by the Google cloud do not have strong requirements for precision frequency or time stamps in sub-second agreement with the time scale in TF.460.

2012: THE LEAP SECOND THAT DID HAPPEN

The leap second of 2012-06-30 occurred in a context of chaos. Even when it had passed the whole story was difficult to piece together. Media coverage of the aftermath was plentiful, but few recognized that there were five different contributors to the chaos, three of which were due to the leap second.

---

* ACR Electronics Inc. (2005-12-05), Media Alert issued by John Bell

† ACR Electronics Inc. (2005-12-05), Technical Bulletin issued by Doug Devore

‡ `http://ask.slashdot.org/story/09/01/01/1930202/`

§ `https://lkml.org/lkml/2009/1/5/562`

¶ `http://googleblog.blogspot.com/2011/09/time-technology-and-leaping-seconds.html`

Figure 2. Google is big enough to define and implement its own global time scale

The confusion began more than a day in advance when a derecho, a storm with powerful line winds, blew across the US from Indiana through Virginia.[*] Damage from the storm included loss of power to the Amazon Elastic Compute Cloud data center in northern Virginia, and this disrupted or disabled many internet and web services. The track of the storm (Figure 3) shows that the loss of power occurred very close to 24 hours before the leap second.

In the 24 hours *before* the leap second, while sysadmins were still recovering from the aftermath of the derecho, some Linux servers began to lock up. Three months earlier, on March 15, there had been a brief discussion in the Linux Kernel Mailing List about a possible livelock bug for systems running NTP.[†] John Stultz had posted a patch for the bug and explained the trade-offs. Unfortunately, by the day of the leap second not all kernels had that patch.

On July 3 sysadmin Bron Gondwana of FastMail.fm blogged his experiences of the day before the leap second.[‡] He described that after rebooting, but still before the leap second, some systems experienced the livelock again. He had started a discussion at ServerFault (which grew during the day before the leap and after) that contains a confusing mix of posts with theories about the effects of the leap second.[§] Some of the entries are about this Linux kernel NTP livelock bug, but many are

[*] `http://en.wikipedia.org/wiki/June_2012_North_American_derecho`
[†] `https://lkml.org/lkml/2012/3/15/616`
[‡] `http://blog.fastmail.fm/2012/07/03/a-story-of-leaping-seconds/`
[§] `http://serverfault.com/questions/403732/`

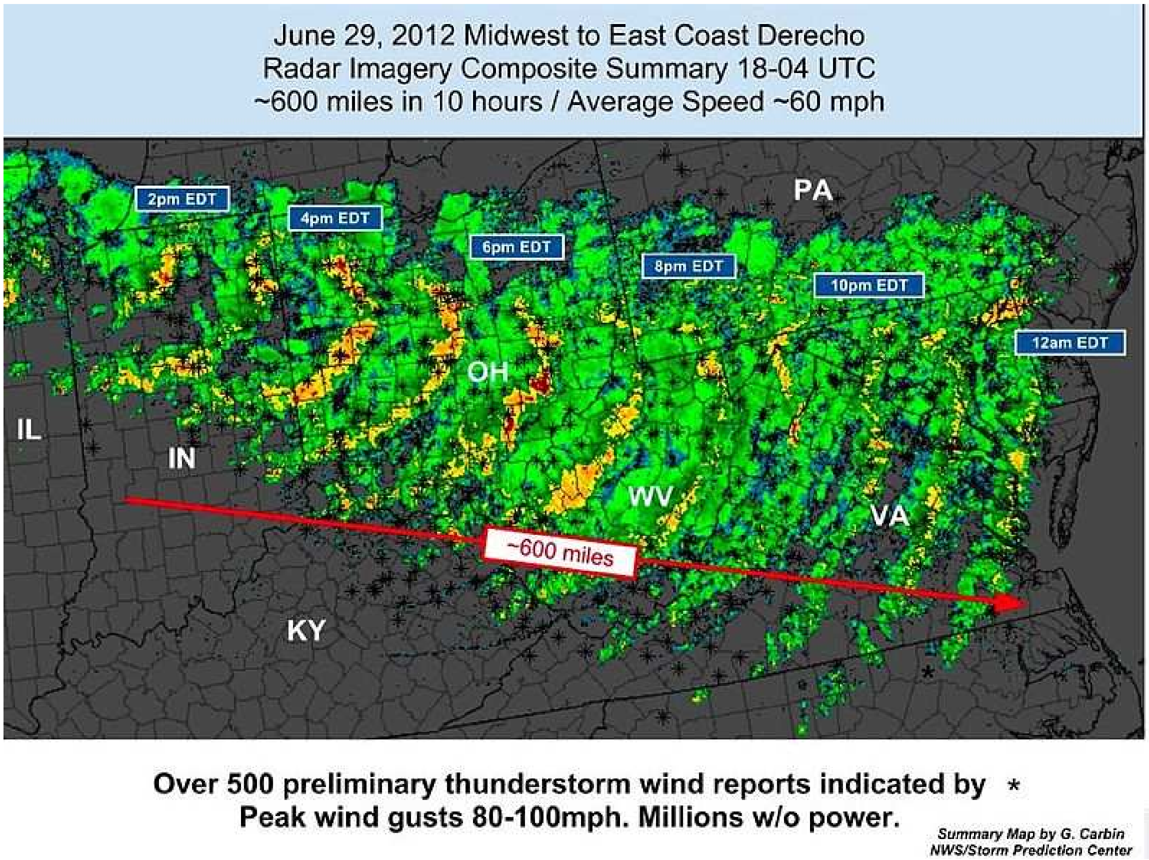

Figure 3. The derecho of 2012-06-29 began a very bad day for system administrators.

about other issues. The likelihood of of a lock depended on the load of a system. Rebooting did not help unless the kernel flags set when NTP is anticipating a leap second were also kept from being reset. That discussion contains many other posts about other Linux kernel problems which did not happen until the leap second, or after.

At the moment of the leap second some systems which had remained unpatched since 2008 (re)experienced the Linux kernel `printk()` deadlock bug. After rebooting, these systems operated with no further problems. In contrast, the remaining Linux kernel problem caused ongoing consternation that did not end for months.

The best known effects of the 2012 leap second came from a previously unrecognized livelock bug in the Linux kernel. After the leap second happened many sysadmins found that their processors were running, but extremely busy doing operations that were not useful. Various blogs indicate that sysadmins at some sites spent hours trying to diagnose their systems before deciding to reboot. News media later reported that air travel around Australia suffered because the Qantas flight reservation system known as Amadeus was unresponsive for hours.

Eventually it became clear that it was possible to restore normal operations simply by issuing the command `date -s "`date`"` to reset the system clock to its current value of the system time. Later in the day John Stulz identified the 2012 Linux kernel futex livelock bug in an apologetic

Linux Kernel Mailing List post*.

The final of the five contributors to the chaos was unrelated to the leap second. The internet still contains many stories with lists of sites that were down around the time of the leap second. Not all of the listed sites were affected by the leap second. In particular the internet servers of The Pirate Bay† and fark‡ had announced that they would be down for scheduled maintenance.

Other systems were affected by the 2012 leap second but never reported in the media. At Lick Observatory on Mt. Hamilton the newest telescope control system for the Automated Planet Finder (APF) noticed the leap second. APF is the only telescope at Lick designed to be fully robotic, and because of that it has the only telescope control system that can be affected by a leap second. As expected, the Unix-based computers applied the leap second on schedule. Also as expected, the Windows-based computers did not apply the leap second. Under normal circumstances this time difference would have been solved by rebooting the affected systems, but in this case the reboot triggered other pre-existing hardware problems. The APF telescope remained offline until after a visit from the vendor.

2012: THE LEAPS THAT DID NOT HAPPEN

At the end of July, the month after the 2012 leap second, various sites again saw all of symptoms of the 2012 leap second. The explanation is easy to see in web pages at Trinity College Dublin School of Mathematics where David Malone monitors many NTP servers.§ Figure 4 shows his plot of the Leap Indicator bits in the NTP protocol. The plot shows that roughly 10% to 20% of the servers in the NTP pool were erroneously announcing an impending leap second.

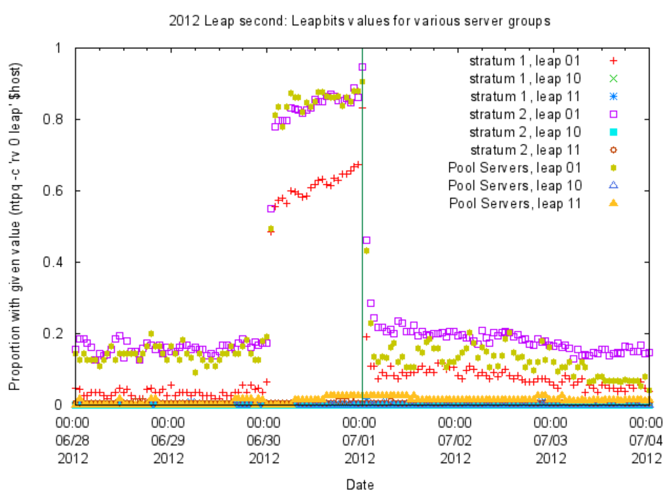

Figure 4. David Malone's plot of NTP Leap Indicator bits

At the beginning of August RedHat (a purveyor of Linux kernels) became aware of the rogue NTP servers and issued a warning to their customers.¶ It seems likely that these false leap sec-

*`https://lkml.org/lkml/2012/7/1/19`

†`http://torrentfreak.com/leap-second-did-not-crash-the-pirate-bay-120701/`

‡`http://www.fark.com/comments/blog284`

§`http://www.maths.tcd.ie/~dwmalone/time/leap2012/#ntpleapflag`

¶`https://access.redhat.com/site/articles/199563` ond announcements prompted many sysadmins not to rely on out-of-the-box installations of the operating system which typically rely on the pool of servers run by the NTP community.*

Nevertheless, on 2013 January 7 operations engineer Chris Brentano posted on the status blog of simple.com that they were unable to handle financial transactions for about two hours on New Year's Eve because their payment processor had been affected by a false leap second.† It is possible that the incorrect NTP servers had not been noticed in previous months because of variations in different versions of NTP servers. Some versions of NTP software only allow leaps at the end of June and December even though TF.460 specifies that they may occur at the end of any month. At UCO/Lick one of our non-critical systems also saw this bogus leap second because its NTP configuration had been left with the vendor-default NTP pool instead of the custom set of NTP servers we maintain.

The biggest surprise leap of 2012 happened on November 19. Van Wolfe was the first of many sysadmins who reported to the North American Network Operators' Group that their systems had jumped from year 2012 back to year 2000.‡ After analysis it became clear that the incorrect date came from two of the NTP servers run by the US Naval Observatory. The discussion of this NTP event at USNO continued on the Outages mailing list where some sysadmins to pondered whether a problem with such an authoritative source might imply some kind of attack.§ The incorrect date triggered authentication issues that prevented logins and caused some sites to reboot all systems.

REAL-TIME SYSTEMS: SKIP NEWHALL'S 2012 LEAP SECOND PARTY

Skip Newhall is an astronomer retired after working in the deep space tracking section of Caltech's Jet Propulsion Laboratory in Pasadena. When a leap second occurs he celebrates with a party.¶ Figure 5 shows his collection of timekeeping gear for his 2012 leap second party.

The behavior of the various devices in rack full of timing equipment gives a hint at the situation faced by real-time operational systems at each leap second. The rack contained 4 GPS receivers, 2 WWVB receivers, a display using IRIG-B from a WWVB receiver, and a display using IRIG-B from a GPS receiver. One of the party guests obtained video of all the displays before and after the leap second.‖ The specifications from the manufacturers of these devices indicated that not all of them would handle the leap second in conformance with TF.460, and they did not. One surprise failure was the display using IRIG-B from the WWVB receivers; its vendor had claimed it would show 23:59:60, but it did not. Another surprise was that one of the GPS receivers displayed a different sequence of time stamps than the other two GPS receivers from the same manufacturer. After the video analysis it is evident that only the two WWVB receivers conformed to TF.460. Subsequent to the leap second, the National Institute of Standards and Technology (NIST) has changed the modulation of the signals from WWVB such that those receivers can no longer achieve precise timing.

CONCLUSION

The disagreement between recommendation TF.460 (which requires leap seconds) and POSIX (which denies their existence) remains a problem. Implementors who are working on systems in-

*http://www.pool.ntp.org/en/

†http://status.simple.com/2013/01/07/nye-postmortem/

‡http://mailman.nanog.org/pipermail/nanog/2012-November/053414.html

§https://puck.nether.net/pipermail/outages/2012-November/004850.html

¶http://www.youtube.com/watch?v=CaOpGrs0x_U

‖Skip Newhall (2012), personal correspondence, Jet Propulsion Laboratory (retired)

Figure 5. Skip Newhall's 2012 leap second party gear at 23:59:60. Only the WWVB receivers conform to TF.460

tended to be robust during leap seconds have to choose which standard they are going to disregard. Purchasers of software and hardware that relies on timing are likely to get a mix where different vendors have chosen to disregard the standards in different ways. The behaviors of a heterogeneous ensemble, or even a homogeneous one, have proved to be unpredictable. Even with the best advance efforts each recent leap second has brought news of unwanted and unexpected surprises for many systems. Reticence to disclose vulnerabilities means that there is no way to know how many more incidents were not reported in the news, or not even recognized as caused by a leap second.

Ongoing evolution of hardware and software suggests that there will continue to be incidents triggered by new and untested aspects of hardware and software. Google has decide to implement the "leap smear" for their non-real-time cloud applications and to employ GPS system time in their Android OS, and those decisions are keeping their systems working despite leap seconds. Sessions at IETF meetings indicate ongoing involvement by hardware and protocol engineers who are developing products that use the TAI-based Precision Time Protocol (PTP, IEEE 1588).[18] Even the ITU-T has chosen to tacitly disregard its 1980 plenary acceptance of UTC by approving the use of PTP in its Recommendation G.8265.1.[*]

I previously opined that getting a good description of leap seconds was like the parable of the blind men and the elephant because the available literature from each authority gives a different set of rules and rationale. Nevertheless, it looks as if developers and system designers have become much more aware of the leap second issues. Perhaps the best example of that was posted in the DevOps Reactions tumblr in early 2013.[†] The motto of that community of system administrators is "Say it with pictures. Describe your feelings about your everyday sysadmin interactions." They

---

[*] `http://www.itu.int/rec/T-REC-G.8265.1/en`

[†] `http://devopsreactions.tumblr.com/post/38053375865/`

employed Johnny Depp's acting skills (Figure 6) in a blog post to communicate their motto:

Figure 6. "When I realized the leap second problem"

Increasing awareness of things that happen during leap seconds is a good thing if it helps to avoid the dire predictions that Commission 31 gave to the IAU in 1970. Nevertheless, the DevOps Reactions blog may also be a sign of another cultural process. It looks like the ongoing effects of leap seconds, and the ongoing failure of all the concerned international agencies to bring closure to the saga of reconsidering leap seconds in the radio broadcast time signals, are giving leap seconds the reputation of internet meme.

REFERENCES

[1] F. Zagar, "Time.," *Transactions of the International Astronomical Union, Series A*, Vol. 14, 1970, pp. 343–355.

[2] CCIR, "Avis 460, Émissions de Fréquences Étalon et de Signaux Horaires," *XII^{e} Assemblée Plénière*, Vol. III, 1970, p. 227.

[3] D. H. Sadler, "Mean Solar Time on the Meridian of Greenwich," *Quarterly Journal of the Royal Astronomical Society*, Vol. 19, Sept. 1978, p. 290.

[4] G. M. R. Winkler, "Time (L'Heure)," *Transactions of the International Astronomical Union, Series B*, Vol. 14, 1971, p. 193.

[5] G. A. Willkins, "Joint Meeting on Time Scales," *Transactions of the International Astronomical Union, Series B*, Vol. 14, 1971, p. 198.

[6] J. Terrien, "Fundamental Metrology," *Metrologia*, Vol. 8, No. 3, 1972, p. 99.

[7] CCIR, "Annex 1/7," *Recommendations and Reports of the CCIR (XIVth Plenary Assembly)*, Vol. XIII, 1978, p. 133.

[8] CCIR, "Minutes of the Eighth Plenary Session," *Recommendations and Reports of the CCIR (XIVth Plenary Assembly)*, Vol. XIII, 1978, p. 65.

[9] International Telecommunication Union, *Final Acts of the World Administrative Radio Conference*. 1980.

[10] CCIR, "Introduction by the Chairman, Study Group 7," *Recommendations and Reports of the CCIR (XVth Plenary Assembly)*, Vol. VII, 1982, p. XVII.

[11] Institute of Electrical and Electronics Engineers, *IEEE Standard Portable Operating System Interface for Computer Environments, Std 1003.1-1988*. 1988.

[12] American National Standards Institute, *American National Standard for Information Systems, Programming language C, ANSI X3.159-1989*. 1989.

[13] Institute of Electrical and Electronics Engineers, *IEEE Standard for Information Technology - Portable Operating System Interfaces (POSIX(R)) - Part 1: System Application Program Interface (API) - Amendment 1: Realtime Extension [C language] Std 1003.1b-1993*. 1993.

[14] The Austin Common Standards Revision Group, *Single UNIX Specification, Version 2*. The Open Group, 1997.

[15] International Organization for Standardization, *Programming languages – C, ISO/IEC 9899:1999*. 1999.

[16] The Austin Common Standards Revision Group, *Single UNIX Specification, Version 3, 2004 Edition*. The Open Group, 2004.

[17] The Austin Common Standards Revision Group, *Single UNIX Specification, Version 4, 2010 Edition*. The Open Group, 2010.

[18] Institute of Electrical and Electronics Engineers, *IEEE 1588-2008: Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*. 2008.