The ISP Column
A column on various things Internet

November 2018
Geoff Huston

What's the Time?

Computers have always had clocks. Well maybe not clocks as you might think, but digital computers have always had oscillators, and if you hook the oscillator to a simple counter then you have a clock. The clock is not just there to tell the time, although it can do that, nor is it there just to record the time when data files are created or modified, though it does that too. Knowing the time is important to many functions, and one of those is security. For example, the instruments we use to describe public key credentials, X.509 certificates, come equipped with validity times, describing the time when a certificate should be used and the time of its expiry. Applications that use these X.509 public key certificates need to compare the times in the certificate to their local clock time as part of the checks to determine whether or not to trust the certificate. One of the assumptions behind this certificate system is that every computer's clock tells the same time. But do they?

This quest for high precision in our clocks is one of humanity's longest quests. Much of the story of precision in the measurement of time is about relating the measurements of the observable universe to the specification of the duration, giving us progressively refined understanding of celestial mechanics. The use of the sexagesimal numeral system for second and minutes dates back to the ancient Sumerians in the 3rd millennium BCE, and this structure ot the measurement of time was passed to the Babylonians, the Greeks and then the ancient Romans.

Mechanical clocks were progressively refined over the centuries, but it was the more practical matter of geolocation of ships, and in particular the calculation of longitude, that provided a significant impetus for vastly improved precision in chronometers in the eighteenth century. John Harrison's H4 marine chronometer, designed in 1761, drifted by 216 seconds in 81 days at sea (or 2.6 seconds per day). Mechanical chronometers continued to be refined, and they kept time for more than a century until the picture changed with the introduction of the quartz crystal oscillator in the early twentieth century.

Quartz oscillators are very stable, and as long as the quartz crystal is not subject to high levels of temperature variation it will operate with a clock drift of less than half a second per day. Far higher precision can be obtained by using a bank of oscillators and also using temperature compensation in the digital counter to obtain a level of accuracy where the clock drift is contained to less than 15ms per day. The United States National Bureau of Standards based the time standard of the United States on quartz clocks from late 1929 until the 1960s.

The search for ever higher precision lead to the use of the oscillation of atoms. The caesium-133 atom oscillates at a level of stability that is within 2 nanoseconds per day, and atomic clocks not only provide us with the reference time measurement but also now also define time itself.

The next question is how to place this highly accurate measurement of a unit of duration of time into the context of the period of the earth's rotation about its axis.

It would be reasonable to expect that the time is just the time, but that is not the case. The Universal Time reference standard has several versions, but two standards are of particular interest to timekeeping.

UT1 is the principal form of Universal Time. Although conceptually it is Mean Solar Time (the average of the period when the sun is at precisely the same elevation in successive days) at 0° longitude, arbitrarily precise measurements of the mean solar time are challenging. UT1 is proportional to the rotation angle of the Earth with respect to a collection of distant quasars, as specified by the International Celestial Reference Frame (ICRF), using long baseline interferometry. This time specification is accurate to around 15 microseconds.

UTC (Coordinated Universal Time) is an atomic timescale that is related to UT1, but at a far greater level of precision. It is the international standard on which civil time is based. UTC ticks in units of "Standard International" (SI) seconds, in step with International Atomic Time (TAI). TAI is based on the definition of 1 second as 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. Because UTC is based on TAI, UTC is not directly related to the rotation of the earth about its axis, but UTC dates still have 86,400 SI seconds per day. UTC is kept within 0.9 seconds of UT1 by the introduction of occasional leap seconds into the UTC day counter.

Many computers use the Network Time Protocol (NTP), or a closely related variant, to maintain the accuracy of the internal clock. NTP uses UTC as the reference clock standard, implying that, like UTC itself, NTP has to incorporate leap second adjustments from time to time.

NTP is an "absolute" time protocol, so that local time zones—and conversion of the absolute time to a calendar date and time with reference to a particular location on the Earth's surface—are not an intrinsic part of the NTP protocol. This conversion from UTC to the wall-clock time, namely the local date and time, is left to the local host.

NTP is remarkably effective, and an NTP client should be able to maintain the local clock with an accuracy of 10ms from the reference clock standard in most cases.

Measuring the Time

At APNIC Labs we use scripted online advertisements to measure various aspects of the Internet. These measurements include the deployment of IPv6 capability, the use of DNSSEC-validating DNS resolvers. We can also use this measurement to report on the variance in clock times.

The measurement process is relatively simple. The script running on the client browser collects the local clock time and passes this value back to the measurement server as part of the script's function. The collected user time can be compared to the server's NTP time.

This is not a very high precision time measurement. There is a variable amount of time between the exact time when the end user program reads the value of the local time of day clock and the time when this information is received by the server. This delay includes the client-side processing delay, the network propagation delay and the server-side processing delay. At best we should admit a level of uncertainty of ± 0.5 second in this form of measurement, and more practically we may want to consider an uncertainty level of ±1 second with this data. But in terms of the overall objective of this measurement exercise this is sufficient for our purpose.

Results

The results of gathering data across a little over 80 days in the third quarter of 2018 reveal that 65,384,331 samples were within one second of the server's reference time. Some 21,524,620 samples were more than 1 second fast, and 115,551,970 samples were more than 1 second slow.

The distribution of relative times in units of days is shown in Figure 1. The scale of the distribution is logarithmic, to show the smaller numbers in the extreme day, but it is evident that a visible number of clients have clocks that are up to one year slower than NTP time. There are clients with local clocks that are up to 2 months ahead of NTP time, and sporadic instances of clients with local clocks that are set to a date that is more than 2 months of NTP time.

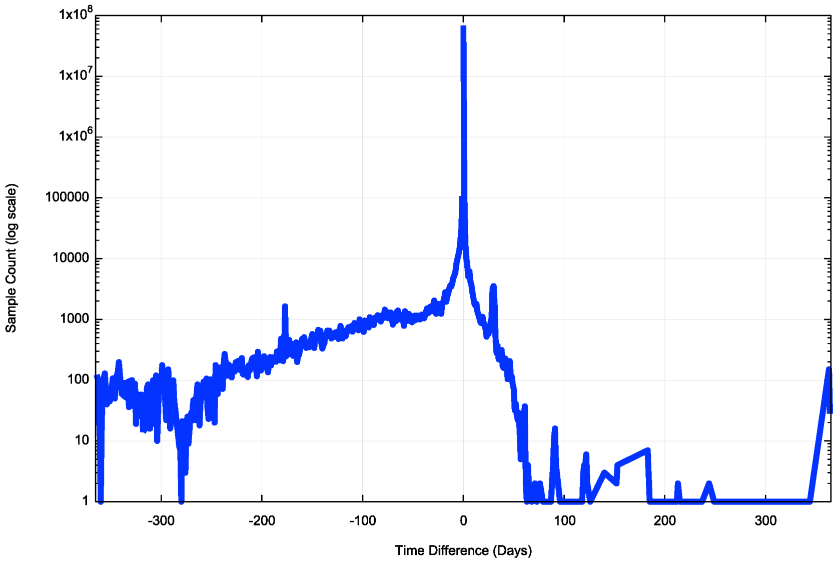

Figure 1 – Distribution of local clock slew

The cumulative distribution of these measurements is shown in Figure 2. This is not a very informative distribution as 92.7% of all samples are within 60 seconds of the NTP reference time.

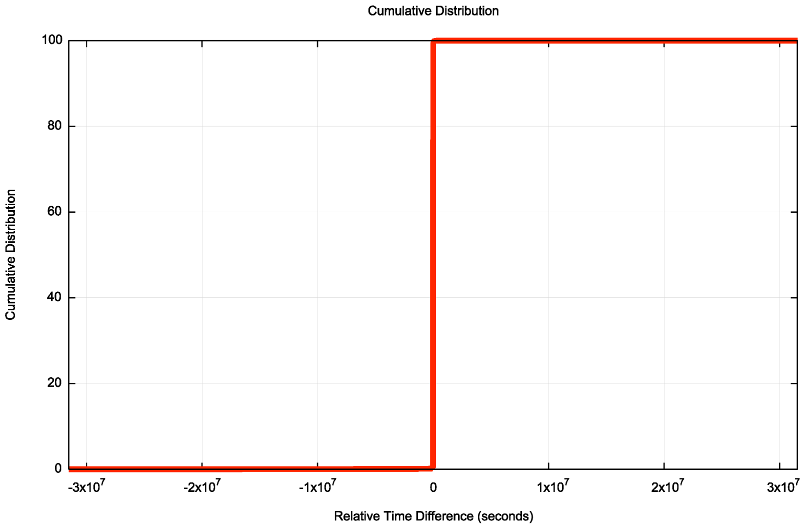

Figure 2 – Cumulative Distribution of local clock slew

We can look at the extremes of this curve, shown in Figures 3a and 3b.

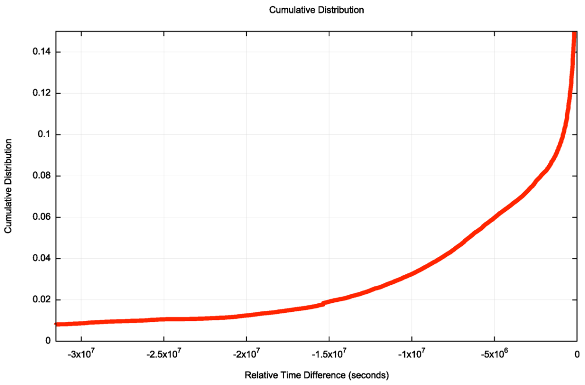

Figure 3a – Negative Clock Slew Distribution

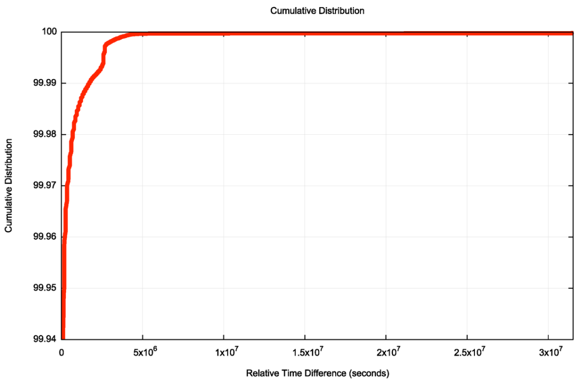

Figure 3b – Positive Clock Slew Distribution

Now let's look at the clients who have clocks that are faster than the server's reference time, and while the distribution is measured in units of seconds, let's use a unit of 1 day on the x axis. What is clear is that of the 0.15% of clients who use local clocks that are ahead of the reference time by one full day or more, the majority of these local clocks are out by multiple of days (Figure 4). These clients appear to have accurate local clocks, yet the mapping of the clock signal to a date is incorrect by a multiple of units of days. There is also a smaller set of clients whose clocks are out by a number of hours (Figure 5). Overall, some 0.7% of clients have local clock settings that are forward of the reference time by one hour or more.

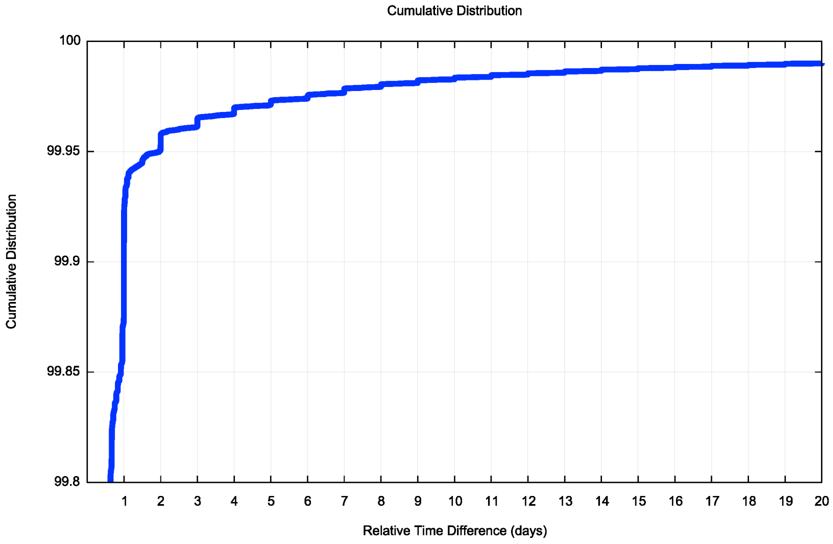

Figure 4 – Positive Clock Slew Distribution (Days)

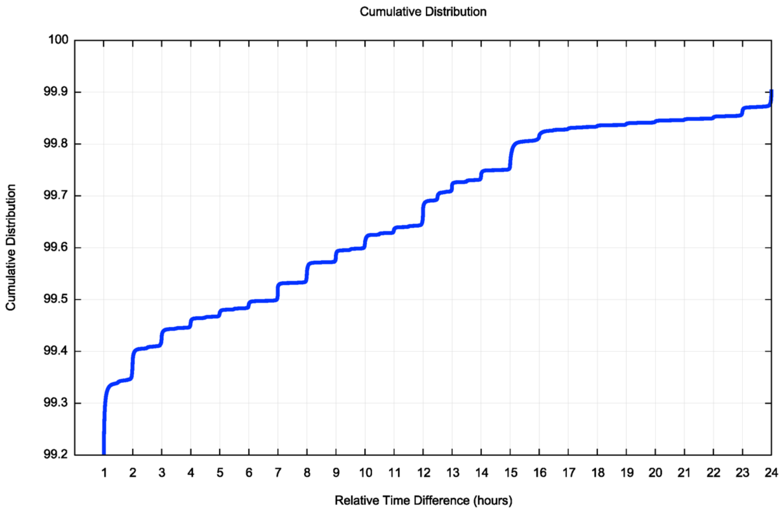

Figure 5 – Positive Clock Slew Distribution (Hours)

What about clients that have local clocks that lag the reference time? Figure 6 shows the cumulative distribution of clients whose clocks lag the reference time by 3 hours or more. While there is some evidence in this figure of clients whose local clock lags the reference standard by exact units of an integral number of days, this is by no means as pronounced as the distribution of clients with faster clocks. It may be larger population of clients with slower clocks (slightly more than twice as many) includes a set of clients whose calendar settings are incorrect and a second set of clients whose local clocks have simply drifted away from reference time, and this time drift can be as large as days and months! However, when we look more closely at just the 24-hour lag period, the quantization of lag into whole units of hours is more pronounced (Figure 7). Some 0.5% of clients have a local clock setting that is exactly 1 hour behind the reference time, and 0.05% of clients have a local clock that is exactly 12 hours behind the reference time.

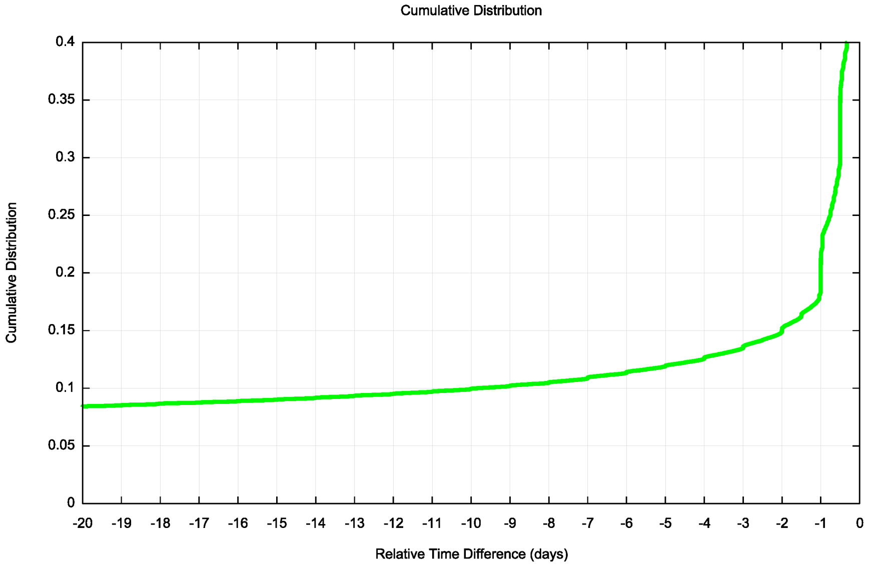

Figure 6 – Negative Clock Slew Distribution (Days)

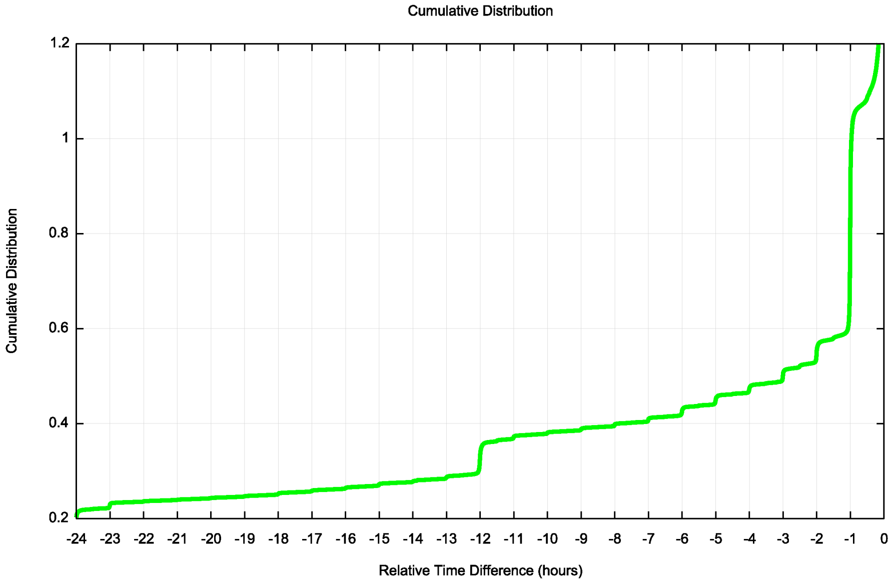

Figure 7 – Negative Clock Slew Distribution (Hours)

This points to an interesting outcome. It is not only clock drift that is the major reason why local clocks are more than a few seconds difference to the reference time, but it is the local time configuration that incorrectly maps what looks to be a stable and accurate local second counter into a local calendar time.

We can see this by looking at the spread of time differences that exclude the major peak point within 1 hour of zero time difference (Figure 8). It is clear that there is a pronounced peak of time differences at the hour chime mark and a secondary peak at the half hour chime mark. It is also evident from this figure that there is a greater propensity to for a local clock to drift to a slower time than a faster time.

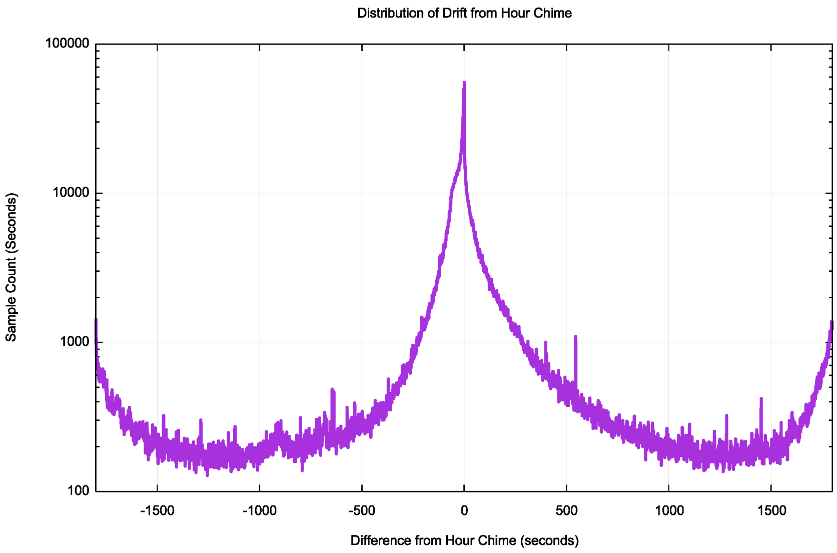

Figure 8 – Clock Slew against Hour Chime (Seconds)

What can we learn from the clock drift data in the period close to exact synchronisation? The distribution of this central hour is shown in Figure 9. Again, there is a greater count of clients with slower clocks than faster clocks.

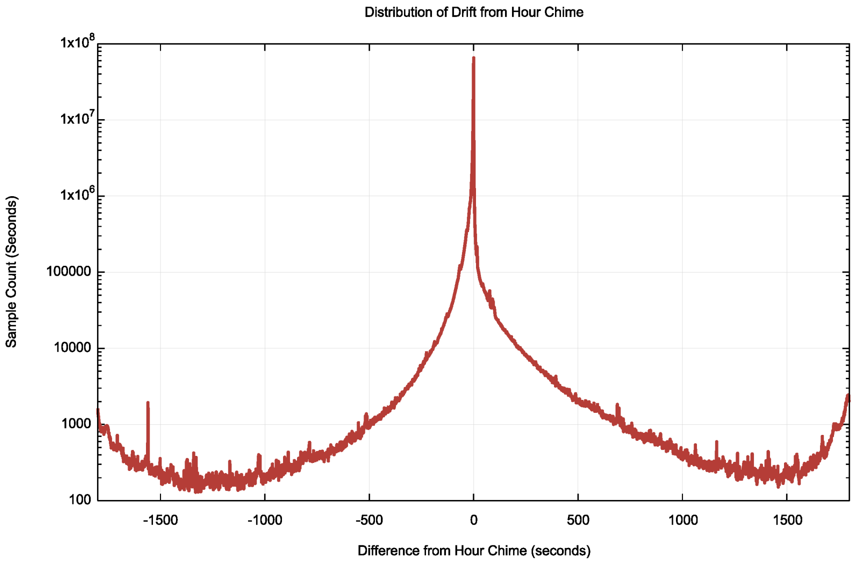

Figure 9 – Clock Slew against First Hour Chime (Seconds)

Summary

Most of the Internet runs to the same clock, but what is "same" admits to a rather broad approximation. Some 53% of clocks on Internet connected devices run within 2 seconds of a server's reference clock, which is close to the limit of accuracy within this experimental method. Of the remainder of the set of tested devices, some 38% of clocks were slow and 9% were fast.

If we broaden this window of approximation of time synchronisation to 60 seconds, then 92% of local clocks are within this 120 second range, and while 4% run fast, and 4% run slow.

If we take an even broader view of time synchronisation, and use a window of ±1 hour, then 98.27% of clients are within 1 hour of the server's reference time, while 0.89% of clients run slower than 1 hour and 0.84% of client run faster than 1 hour.

The drift of time in these local clocks where there is significant variance is not uniform. The strong quantisation of the clock drift into units of hours tends to suggest that a major component of this clock slew is not the drift of the local oscillator or dropping of clock ticks in the time management subsystem, but some form of misconfiguration of the local date calculation. The second counter appears to be quite stable, but the local date calculation is off.

What can we say about time and the Internet? The safest assumption is that most systems will be in sync with a UTC reference clock source as long as the definition of "in sync" is a window of 24 hours! If the 'correct' behaviour of an Internet application relies on a tighter level of time convergence with UTC time than this rather large window, then it's likely that a set of clients will fall outside of the application's view of what's acceptable.

Disclaimer

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

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net