Time Traceability for the Finance Sector

Fact Sheet

Version 1.4
14 March 2016

NPL Management Ltd is a company registered in England and Wales No. 2937881
Registered Office: NPL Management Ltd, Hampton Road, Teddington. Middlesex, United Kingdom TW11 0LW

Contents

1. Introduction

This document describes the key concepts that underpin the emerging requirements for accurate time-stamping of transactions in the financial sector. The draft regulatory technical standard for clock synchronisation is summarised, and subsequent topics covered include the international time system and the importance of UTC, the concept of traceability for time measurements, and the terminology used for traceability in time dissemination and time stamping.

Accurate and trusted timing plays a critical role in financial markets, underpinning the time stamping of trades, the synchronisation of computer systems and the measurement of network latency for process optimisation. The rapid expansion of computer-based trading has increased the need for synchronisation of trading systems and traceability to a common reference time scale, to help prevent trading irregularities and to aid forensics. To encourage this, a number of regulatory bodies now require high precision traceable time-stamping to help understand activity across trading venues, working toward a consolidated audit trail.

Time stamps created by different systems or networks can only be compared meaningfully if they are based on the same reference. For time, the global reference is Coordinated Universal Time (UTC): the time scale that underpins GPS, broadcast time signals and all other precise time services. UTC is generated by the International Bureau of Weights and Measures (BIPM) through an international collaboration involving around 75 timing institutes. Each of these institutes maintains a physical representation of UTC, called generically UTC(*k*), which can act as the reference for national or regional time dissemination services.

Whatever time source and dissemination method are used, they must provide traceability back to UTC. Traceability requires a continuous chain of comparisons with known uncertainties, all of which must be documented. The GPS satellite signals alone do not readily provide traceability to UTC, but users can demonstrate traceability by obtaining GPS monitoring bulletins from one of the regional UTC(*k*) timing centres.

There are two other requirements for demonstrating time traceability. One is for the timing equipment to be calibrated, so that its unknown internal delays do not bias its time output. The other is for the equipment to be monitored continuously, so that any fault or anomalous behaviour can be detected and the time output not used until the equipment is working correctly again. The calibration evidence and monitoring results should be archived so that the status of the timing equipment at any point in time can be verified at a later date.

2. Technical standards for business clock synchronisation

The European Securities and Markets Authority (ESMA) has issued a series of draft technical standards in support of the MiFID (Markets in Financial Instruments Directive) II regulations, which are awaiting approval by the European Commission. One of the supporting Regulatory Technical Standards, RTS 25, deals with clock synchronisation. Its technical regulations are described under four articles, which are summarised below.

Article 1 states that business clocks that give the timestamp for any reportable event should be synchronised to Coordinated Universal Time (UTC), using either a link to one of the timing centres maintaining a UTC(*k*) realisation of UTC or the timed signals disseminated by GPS or another satellite system. The concept of traceability is mentioned earlier in RTS 25.

Article 2 describes the level of accuracy (maximum divergence from UTC) that should be achieved by the operators of trading venues, taking into account the gateway to gateway latency of their trading systems. These are given in table 1.

Latency time	**Max. divergence from UTC**	**Timestamp granularity**
>1 millisecond	1 millisecond	1 millisecond or better
=< 1 millisecond	100 microseconds	1 microsecond or better

Table 1: Level of accuracy for operators of trading venues.

Voice trading systems or those needing human intervention need only ensure 1 second maximum divergence from UTC, and record times with a 1 second granularity.

Article 3 defines the level of accuracy that apply to members or participants of trading venues, as summarised in table 2.

Type of trading activity	**Max. divergence from UTC**	**Timestamp granularity**
High frequency algorithmic	100 microseconds	1 microsecond or better
Voice trading systems	1 second	1 second or better
Human intervention; non-algorithmic	1 second	1 second or better
Concluding negotiated transactions	1 second	1 second or better
Other	1 millisecond	1 millisecond or better

Table 2: Level of accuracy for members or participants of a trading venue.

Article 4 specifies the need to demonstrate traceability to UTC by documenting system design, functioning and specifications, and to identify the exact point at which a timestamp is applied. The traceability system should be reviewed at least once a year to ensure compliance with the regulations.

3. Global timekeeping

Accurate timekeeping worldwide is based on a single reference time scale, known as Coordinated Universal Time and abbreviated in all languages as UTC. The International Bureau of Weights and Measures (BIPM), based in Paris, computes UTC in monthly blocks by combining data from around 400 continuously-operating atomic clocks located in national timing centres. The duration of the UTC second is fine-tuned using measurements from a small number of primary frequency standards to ensure that it remains as close as possible to the standard unit of time, the *second*, defined in the International System of Units (SI).

UTC is computed monthly, so does not exist in real-time. Each institute contributing clock data to the BIPM maintains its own physical realisation of UTC, such as UTC(NIST) in the USA and UTC(NPL) in the UK, which are known collectively as the UTC(*k*) time scales. These national time scales are adjusted so that they remain close to UTC, usually within 1 microsecond, and in some cases the difference is kept below 10 nanoseconds. They are traceable to UTC, and serve as the reference standards for all accurate time measurements globally.

UTC is based on atomic clocks, giving it great stability and accuracy. However, it slowly deviates from time based on the Earth's rotation, which fluctuates unpredictable over time and experiences a long-term slowing due to friction caused by the tides. The increasing difference between UTC and Earth rotation time (UT1, which can be thought of as a more precisely-defined version of Greenwich Mean Time) is occasionally corrected by a 1-second adjustment in UTC known as a leap second. Since the mechanism was introduced in 1972 there have been 26 leap seconds. They have so far all been positive, so the final minute of the day has 61 seconds rather than 60, and applied at the end of either 30 June or 31 December.

4. Sources of time

Time is readily available from a broad range of sources and over many different distribution protocols, with considerable variations in the resolution, stability and accuracy that they offer the user. All disseminated time signals and services are derived from UTC, but in practice it can be difficult or impossible to validate the entire distribution chain well enough to calculate the uncertainty of every link and establish traceability.

Commonly used sources of time include:

1. **Free-running clock.** A clock on its own, even a caesium atomic clock, is not a viable source of traceable time-of-day as its time and frequency offsets from UTC are unknown and will change over time. All clocks are affected by a range of noise processes and will tend to diverge from UTC even if synchronised (set to the same time) and syntonised (set to the same rate) initially. Averaging an ensemble of several clocks does not help greatly as the ensemble will itself drift even if the individual clocks do not have similar behaviours.
2. **GPS.** The Global Positioning System and other navigation satellite systems are great sources of accurate time, and in a typical installation a GPS-disciplined oscillator will disseminate time across a local-area network using NTP or PTP (discussed below). However, such solutions need to be implemented with care[1]. Potential sources of error include multipath reflections of the

[1] Global Navigation Space Systems: reliance and vulnerabilities, The Royal Academy of Engineering, Mar 2011: http://www.raeng.org.uk/publications/reports/global-navigation-space-systems

satellite signals, space weather events[2], RF interference at the antenna, and uncalibrated delays in antenna cables or receiver hardware. A GPS-disciplined oscillator feeding a network time distribution must be monitored or compared continuously with an independent time source to verify that it remains locked to the satellite signals.

3. **NTP servers.** The Network Time Protocol (NTP) is widely used to disseminate time over the internet and large numbers of servers can be found online. Servers in unknown locations should be avoided as many are based on GPS-disciplined oscillators and their performance may be affected by local factors such as interference and multipath effects. Most UTC(*k*) labs operate NTP servers that are monitored and synchronised to their time scales. These servers can in principle deliver traceability to UTC, particularly if an NTP authentication method such as MD5 is employed. However, the NTP protocol can only provide synchronisation over wide-area networks to a few tens of milliseconds and will not meet the MiFID II requirements.
4. **Standard-frequency and time signals.** A number of countries operate radio signals that provide access to time based on a UTC(*k*) time scale. In general, though, the accuracy of the signals varies from 10s of milliseconds down to 10s of microseconds, depending on the form of the modulation and on variations in the signal propagation, and they are rarely used for network synchronisation.
5. **UTC(*k*) delivery over fibre**. NPL*Time*® and similar services disseminate UTC-traceable time over managed fibre links using the Precision Time Protocol (PTP version 2, defined in the standard IEEE 1588-2008). PTP is a dissemination method originally developed for local area networks and is capable of achieving accuracy better than 100 nanoseconds over stable and symmetric links. PTP can deliver synchronisation over longer distances using telecoms fibre networks, employing dedicated channels or PTP-compatible switches to maintain accuracy and offering MiFID II compliance at the ingress point of the customer distribution system. As with NTP, the latency to each end point is continuously measured by the protocol (assuming out and back symmetry) and the offset corrected. The major risk to this type of service is of the fibre link being severed, due to roadworks for example, but local holdover mechanisms can be put in place to maintain service provision whilst the repair is being effected.

In summary, it is not enough to have an accurate time source feeding a distribution network: it is critical to understand, manage and document how time is delivered, distributed and consumed. At each step in the process, knowledge of the traceability of the time signal to UTC is essential to ensure regulatory compliance at the time-stamp. Achieving MiFID II compliancy at the time-stamp requires a complete understanding of the traceability chain and the uncertainty of every link, all the way from UTC to the time-stamp, together with an archive of the evidence needed to demonstrate this knowledge.

5. Traceability and uncertainty

The concept of *traceability* for measurements of time-of-day refers to a continuous chain of comparisons extending from a time comparison, time stamp or clock synchronisation performed by a user, back through the distribution to one of the UTC(*k*) time scales, and so to the reference time scale UTC. Each comparison in the chain will be carried out with some inherent level of inaccuracy, which can be quantified as the *uncertainty* of the measurement. The uncertainties of each link in the dissemination solution can then be combined to give the total uncertainty of the time signal at the point where it is available to the user.

[2] Extreme space weather: impacts on engineered systems and infrastructure, The Royal Academy of Engineering, Feb 2013:
http://www.raeng.org.uk/publications/reports/space-weather-full-report

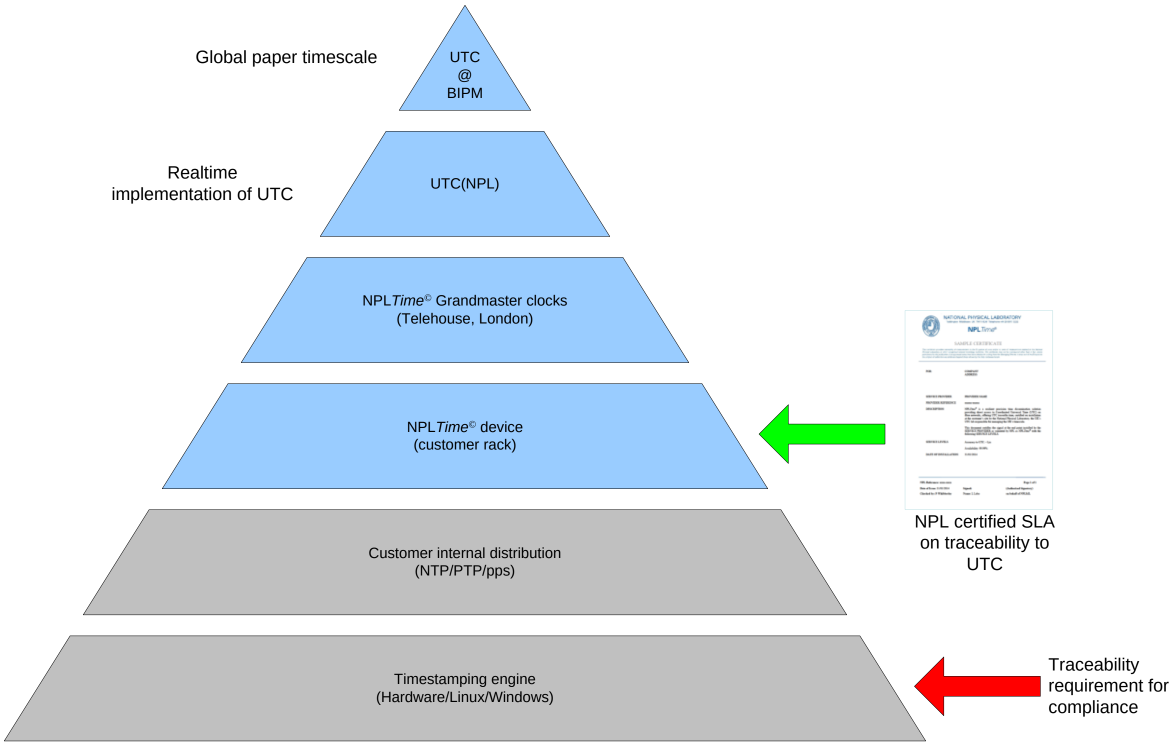

Figure 1: Traceability pyramid detailing the comparisons linking a timestamp to UTC, taking the NPLTime service as an example.

Understanding the dissemination chain and calculating the uncertainty is an essential requirement for demonstrating traceability. Each link must be documented, along with the associated uncertainty evaluation, and the total uncertainty of the timestamp or time output must be determined and recorded.

For example, an NTP service operated by a UTC(*k*) institute feeding an NTP internal distribution, culminating in a Windows based time-stamping engine, may have the following *uncertainty budget*:

Link in distribution chain	Standard uncertainty
Timestamping engine	0.1 s
NTP distribution over LAN	0.01 s
NTP distribution over the internet	0.05 s
Sync of NTP server to UTC(*k*)	0.000001 s
Total uncertainty	0.11 s

The combined standard uncertainty[3] of the traceability chain between the user and UTC can generally be calculated as the quadrature sum of the individual uncertainties of each link:

$$u_c = \sqrt{a^2 + b^2 + c^2 + \cdots}$$

$$= \sqrt{(0.1)^2 + (0.01)^2 + (0.05)^2 + \ (0.000001)^2}\ s$$

$$= 0.11\text{s}$$

Note that to determine a combined uncertainty value, the uncertainty of each link in the traceability chain has to be measured or otherwise determined, for example in a calibration carried out by an accredited institute.

In addition, it is not sufficient simply to evaluate the uncertainty budget at one point in time. There must be continuous monitoring at key points of the distribution to demonstrate that the traceability chain remains valid and the time available at each point is within the calculated uncertainty at all times when measurements are being taken. The monitoring records must therefore be stored so that they can be checked at some later time to confirm that the equipment was operating correctly.

6. Obtaining traceability using GPS

GPS receivers have become widely used as reference time standards, providing synchronisation of devices across local area networks. Most of these devices are GPS disciplined oscillators (GPSDOs). Their 'self-adjusting' behaviour and ease of use make them an attractive choice for many applications. However, it is very difficult to demonstrate direct traceability to UTC using a GPSDO, and they should be installed and used with care.

A GPSDO contains an internal oscillator, usually a quartz crystal oscillator or a rubidium atomic frequency standard, that generates signals with good short-term frequency stability. Long-term accuracy is obtained by steering the oscillator to the time scale broadcast by the GPS satellites, known as GPS time, which derived from the realisation of UTC at the United States Naval Observatory, UTC(USNO). GPS time does not implement leap seconds so has an integer-second offset from UTC(USNO), but the current offset and forthcoming changes are broadcast within the GPS navigation message. A GPSDO applies the offset to its time output so that it provides a representation of UTC rather than GPS time. This steering procedure enables a GPSDO to deliver a high level of performance, and to maintain that performance indefinitely.

[3] A Beginner's Guide to Uncertainty of Measurement, Stephanie Bell, NPL, Mar 2001: https://www.wmo.int/pages/prog/gcos/documents/gruanmanuals/UK_NPL/mgpg11.pdf

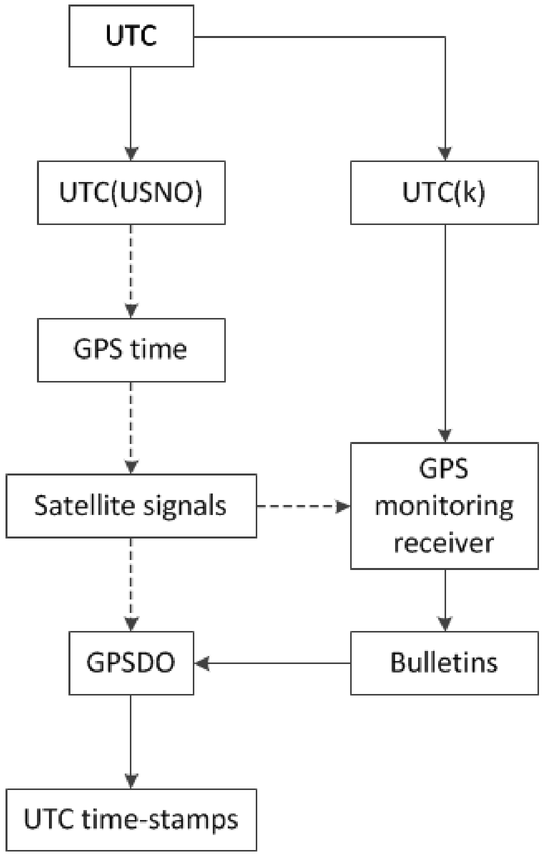

Figure 2: Chain of comparisons from UTC to the time-stamps generated by a GPS disciplined oscillator, and (on the right) the use of bulletins of GPS monitoring results from a UTC(k) institute to demonstrate traceability to UTC. The dashed lines indicate links where it can be difficult or impossible to demonstrate traceability. Note that although the pulses of GPS time are aligned closely with those of UTC, there is an integer-second difference between them (contained in the satellite signals) which changes when a leap second is inserted into UTC.

The distribution chain from UTC to the time output of a GPSDO is shown on the left side of figure 2. The links represented by dashed lines cannot easily be evaluated and assigned an uncertainty by an external user, making direct traceability difficult to establish.

A solution often adopted by calibration laboratories is to subscribe to one of the GPS monitoring bulletins published daily, weekly or monthly by some national measurement institutes. These bulletins give a measured value for the time difference between that institute's UTC(*k*) time scale and GPS time, providing traceability between the satellite signals observed in that region and UTC.

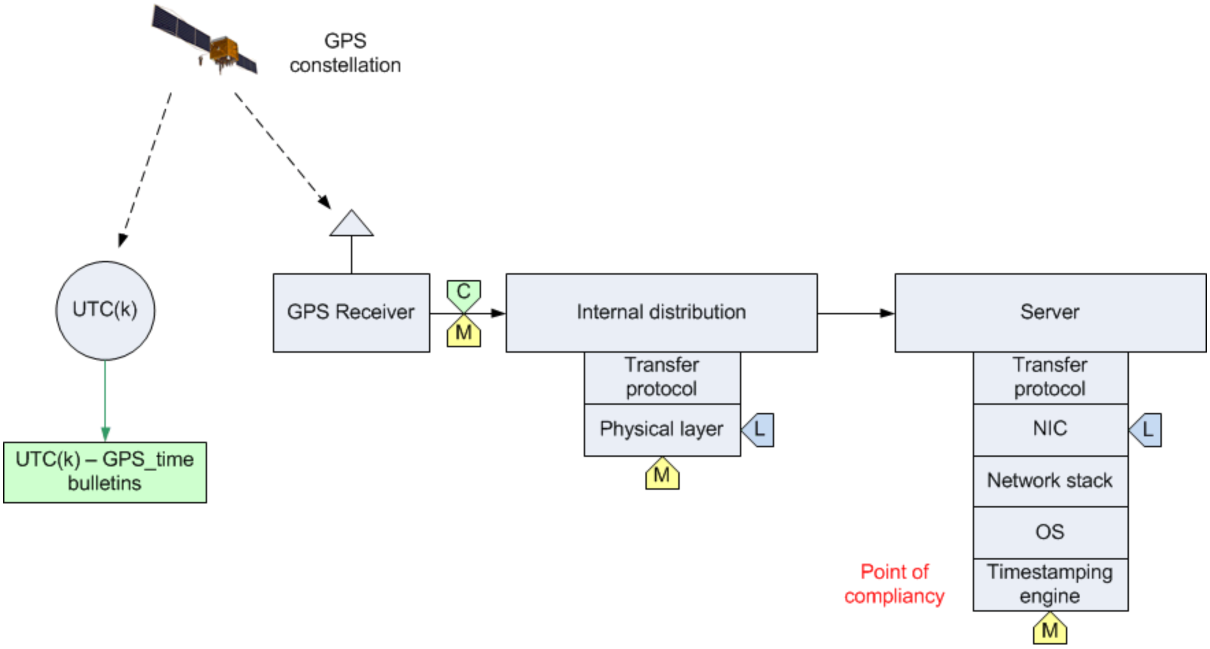

Figure 3: Diagram of a typical GPS-based distribution, indicating the traceability chain. 'C' indicates the point where a calibration may be carried out to determine the offset introduced by the GPS receiver, 'L' indicates points where latency (distribution offsets) can be determined, and 'M' indicates points where continuous monitoring of the time distribution can be carried out.

Even if the principle of traceability through the satellite signals is accepted, there are two other aspects that must be considered before UTC traceability can be claimed for the time output of a GPSDO.

The first requirement is for calibration of the GPSDO. Although the time derived from the satellite signals should be aligned closely with UTC in the long-term, the output signal from the GPSDO may have a constant offset due to uncorrected internal delays, including the delay of the antenna cable. The total delay of the equipment can be determined in a calibration carried out by a suitably accredited body, and the appropriate correction applied to remove the offset. In some situations, for example where the claimed uncertainty of the UTC output is 1 microsecond or more, it may be possible to use the GPS receiver manufacturer's own evaluation of its time offset to satisfy the traceability requirement. However, in most cases a calibration will be needed.

The second requirement is to have some method of monitoring a GPSDO acting as a reference time source to verify that it remains correctly synchronised to the satellite signals. The most effective solution is to compare the output signal from the GPSDO continuously with the similar output from another device. A sudden change in the time difference between the two devices will indicate that there is a problem with one of them, and neither should be used as a reference standard until the fault has been identified and fixed. The second device can be either a free-running clock or a different model of GPS receiver, preferably from a different manufacturer (since two receivers of the same type may misbehave in the same way). If a second GPS receiver is used, it should be a completely separate installation with its own antenna, mounted well away from the antenna of the primary GPSDO to minimise the risk of both receivers being affected in the same way by factors such as multipath reflections or local interference.

If a second standard is not available, an alternative approach is to monitor the operating parameters of the GPSDO and to check for changes or discrepancies. The parameters observed might include the oscillator locking voltage, the number of satellites tracked, or the satellite signal strengths. Whatever method is adopted, the monitoring results should be recorded and archived so that the status of the GPS receiver at any point in time can be checked at a later date.

7. Conclusions

The emerging MiFID II regulations in Europe require time-stamping of automated transactions with an accuracy of 100 microseconds (μs). To have any validity, the time source and dissemination method used must be traceable to UTC, the international reference time scale. In practice, the measurement should relate back through a traceability chain of comparisons with known uncertainties to one of the timing institutes that maintain a physical representation of UTC, or UTC(*k*) time scale. The traceability chain and the uncertainty evaluation must be documented thoroughly.

Traceability is particularly difficult to demonstrate when using a GPS receiver as the time source, but the GPS monitoring bulletins issued by one of the regional UTC(*k*) timing centres provide an alternative route.

Whatever timing equipment is used, there are two other requirements for demonstrating traceability to UTC. The first is that the timing equipment should be calibrated to determine and if necessary correct the error in its time output due to its internal delays. The other is that the equipment should be monitored continuously to detect any fault that might affect its time output. Monitoring and calibration records must be kept to provide evidence of correct operation.

Appendix 1: Terminology of time measurement

Brief explanations of some of the more common terms used in time measurement are given here.

A**ccuracy** of a measurement represents its closeness of agreement with the true value, and can be quantified by assigning an uncertainty to it. For time measurements, the 'true value' is provided by the international reference time scale UTC.

Precision is the level of agreement between repeated measurements of a quantity, and can often be calculated using a statistical tool such as a standard deviation or a variance.

Resolution is the smallest difference that can be observed in a measurement, and is dependent on the properties of the measuring device.

Stability (or equivalently **instability**) is a measure of the fluctuations in a measured quantity over a given time interval. The dominant noise processes observed in a clock or time distribution system depend on the measurement time, and the stability is quantified using statistics such as the Allan deviation that are a function of the sampling period.

Traceability of the time given by a clock requires an unbroken chain of measurements that link it back to the agreed reference standard, which for accurate timekeeping worldwide is UTC. A timestamp can only be considered accurate if its traceability back to a source of UTC has been established.

Uncertainty of a measurement is a numerical assessment of the level of doubt in the result. The uncertainty can often be separated into two components. The Type A uncertainty, u_A, is a statistical value expressing the random fluctuations in the measurement result, usually determined by taking repeated measurements. The Type B uncertainty, u_B, is an estimation of the systematic bias in the measurement result. If these two components are independent, a total uncertainty u can be calculated as their root-sum-square combination: $u = \sqrt{{u_A}^2 + {u_B}^2}$

Coordinated Universal Time (UTC) is the international reference time scale. It is computed monthly by the International Bureau of Weights and Measures (BIPM) using more than 400 atomic clocks in around 70 timing centres world-wide. These timing centres maintain real-time realisations of UTC, referred to as UTC(*k*), where *k* designates the abbreviation allocated to an institute. The scale interval of UTC is the SI second.

The SI second (s) is the basic unit of time in the International System of Units (SI). It is defined as the duration of a specified number of cycles of the microwave radiation corresponding to a particular transition within caesium 133 atoms. The second can be subdivided into milliseconds (1 s = 1000 ms), microseconds (1 s = 1,000,000 μs) and nanoseconds (1 s = 1,000,000,000 ns).

Appendix 2: Example of NPL monthly GPS monitoring bulletin

Several National Measurement Institutes in Europe monitor continuously the GPS time broadcast in the GPS satellite signals against their respective UTC(*k*) time scales, and publish the a summary of the measurements in weekly or monthly bulletins. An NPL GPS monthly bulletin is included below as an example.

NATIONAL PHYSICAL LABORATORY
Time and Frequency Services
Teddington, Middlesex, TW11 0LW, United Kingdom

Web site: www.npl.co.uk/science-technology/time-frequency/

NPL GPS Bulletin

No.2016-01 January 2016

MJD	Date	[UTC(NPL) - GPS_time] mod 1s (ns)
57388	2016-01-01	5.0
57389	2016-01-02	4.6
57390	2016-01-03	5.1
57391	2016-01-04	5.0
57392	2016-01-05	5.5
57393	2016-01-06	6.1
57394	2016-01-07	6.8
57395	2016-01-08	7.4
57396	2016-01-09	8.2
57397	2016-01-10	8.3
57398	2016-01-11	8.9
57399	2016-01-12	9.5
57400	2016-01-13	10.0
57401	2016-01-14	9.2
57402	2016-01-15	8.6
57403	2016-01-16	8.1
57404	2016-01-17	6.7
57405	2016-01-18	5.8
57406	2016-01-19	6.1
57407	2016-01-20	5.4

57408	2016-01-21	5.5
57409	2016-01-22	7.0
57410	2016-01-23	6.7
57411	2016-01-24	7.0
57412	2016-01-25	5.8
57413	2016-01-26	4.6
57414	2016-01-27	3.9
57415	2016-01-28	1.1
57416	2016-01-29	-1.6
57417	2016-01-30	-2.5
57418	2016-01-31	-2.6

NOTES:

1. #.# indicates that NPL data are not available.
2. The total 95% confidence interval on each daily value is +/-22ns.
3. Due to leap second insertions in UTC, [UTC(NPL)-GPS_time] div 1s = -17s.
4. UTC(NPL)-GPS_time = [UTC(NPL)-GPS_time] div 1s + [UTC(NPL)-GPS_time] mod 1s.
5. Expressed in words, UTC(NPL)-GPS_time difference = leap seconds + column data.
6. This report has been compiled by GPSMONITOR201.EXE version 2.01.
7. The measurements in this report were taken by Dicom GTR50 GPS time transfer receiver s/no 0807183.
8. The measurements in this report are single-frequency C/A code observations with the ionospheric delay corrected using a P3 combination of the P1 and P2 code measurements.
9. No anomalous GPS measurements were detected during the period covered by this report.