120 NANSOSECOND WORST CASE NTP PERFORMANCE WITH TIMEKEEPER®

FSMLABS TECHNICAL STAFF

TimeKeeper®routinely produces sub-microsecond accuracy from NTP sources. In the graph below, a TimeKeeper client tracks an NTP source to within *140 nanoseconds*. The client is setup with four sources: two independent GPS sources for validation (blue and green), a PTP source (yellow), and an NTP source (orange). The worst case (not average) NTP source accuracy is within 140 nanoseconds and is generally far better than that.

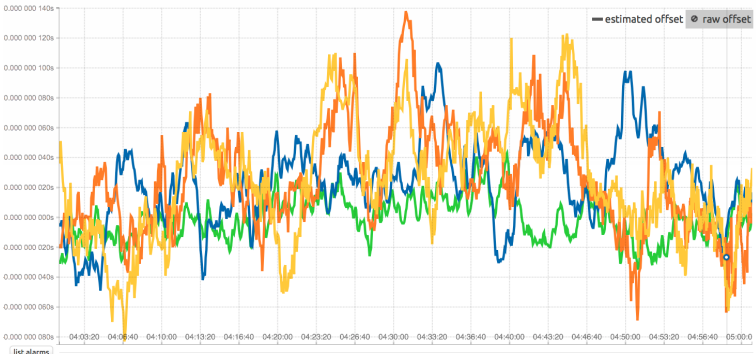

FIGURE 1. Worst case NTP accuracy is better than 140 nanoseconds

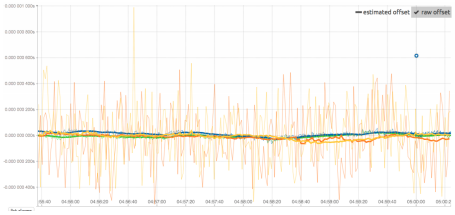

FIGURE 2. Raw and smoothed

This level of accuracy is no accident: not only is TimeKeeper® making aggressive use of hardware timestamping network cards, but it is applying both sophisticated machine learning based filtering and smoothing *and* it is making use of some extensions to the protocol that permit interoperability but significantly reduce jitter[4]. And this NTP feed comes from a TimeKeeper GrandMaster. Network clocks using NTPd to serve NTP will not come close. In the drill down graph to the left, we show a steady clock from NTP with jitter on both raw NTP and PTP (dotted lines) that is smoothed out by TimeKeeper®.

(1) NTP accuracy is excellent, well below 1/2 microsecond (PTP is good too!).

Date: May 11 2018.
Key words and phrases. NTP, PTP, clock synchronization.

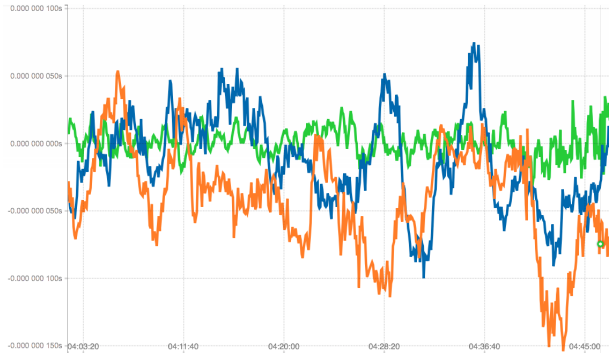

FIGURE 3. GrandMaster graphs of time coming back from Client

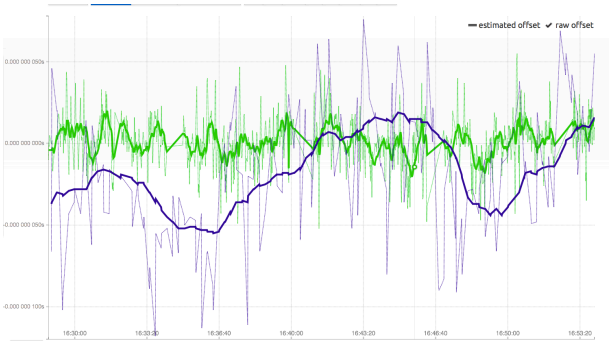

FIGURE 4. GrandMaster graphs of Client reported time versus GM GPS

(2) The accuracy numbers are *meaningful* and *validated* against the two reference GPS sources. In this test, GPS time is not used to correct the NTP time, but to provide a comparison.

To increase confidence, we can look at the monitoring information from the TimeKeeper GM. For figure 3 we did a simple test where the client was asked to be a source, feeding back to the GM, and the GM tested both PTP and NTP against its GPS source. The graph shows that the approximately 150 nanosecond jitter is preserved in this measurement as well. TimeKeeper also monitors the reported time on the client from the server. That time is often much better than the reported accuracy because it results from much deeper analysis of the clock (see figure 4).

Measurement. TimeKeeper incorporates a significant amount of engineering to make its accuracy measurements meaningful. But it is not unusual to see clock synchronization client software that grossly under-reports errors and that undergoes silent failures. For example, a test using NTPd[1], a NTP client often bundled with standard operating systems, reported that the clock was near perfect in a virtual machine that had just been suspended for several seconds. The suspend operation stopped the clock, so when the system restarted the system clock was several seconds behind the actual time, a situation that NTPd ignored. The free software PTP client, PTPd, will often fail to detect network configuration problems that make high quality synchronization impossible but report good sync nevertheless. In fact, the whole EUREX trading network

was inoperable for several hours in 2013 due to a minor failure in a single satellite receiver clock which reported its own accuracy to be nearly perfect.

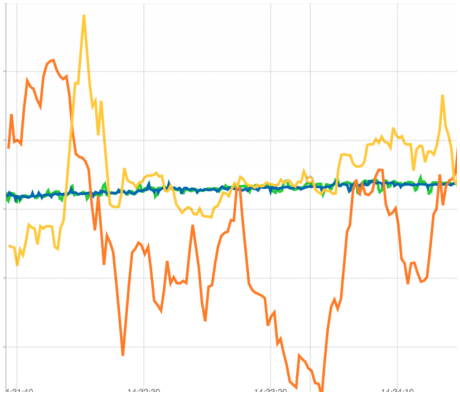

FIGURE 5. Frequency

The unreliable nature of self-reported accuracy in some clock synchronization clients should be no surprise. Not only does estimation of accuracy depend on correctness of clock model and estimations of everything from network delays to the effects of temperature on local oscillators, but it is essentially impossible with a single clock source. Estimating accuracy is complicated: the level of regression test and measurement needed to make sure accuracy is correctly calculated is quite high. Software that comes out of a less disciplined development process may have wildly variable behavior between even minor versions. And fundamentally, it is impossible for any clock synchronization client software, even TimeKeeper, to determine clock accuracy without some reference for comparison. **At best, with a single source, "accuracy" means the accuracy with which client software can track the source.** The client has no way to determine if the single source is anywhere near reference time, whether its own calculations are distorted by network asymetric delays, or whether its local oscillator is even stable! To improve accuracy estimates, TimeKeeper analyzes clocks from several aspects, including frequency (see on the right) and collects additional data such as temperature when available. But reliable verification without multiple clock sources [3, 2] is impossible.

Summary and Warning. Contrary to folk wisdom, it is possible to get very high accuracy from NTP. These test system, however, depends on a optimized configuration. It is possible to improve these numbers by using higher quality hardware and a better switch, but don't expect to duplicate the same numbers on an underpowered virtual machine receiving time over a wide area network from a sub-optimal NTP source. In fact, don't expect to get close to these numbers with Stratum servers that are not running TimeKeeper.

REFERENCES

[1] Bruce Byfield. *A rift in the NTP world*. `https://lwn.net/Articles/713901`.

[2] Cort Dougan and Victor Yodaiken. *US Patent 9348321: Method, time consumer system, and computer program product for maintaining accurate time on an ideal clock*. Austin, TX, 2013.

[3] Cort Dougan and Victor Yodaiken. *US Patent 9671761: Method, time consumer system, and computer program product for maintaining accurate time on an ideal clock*. Austin, TX, 2016.

[4] Cort Dougan and Victor Yodaiken. *US Patent 9756153: Method for improving accuracy in computation of one-way transfer time for network time synchronization*. Austin, TX, May 2012.

Austin Texas.
E-mail address: info@fsmlabs.com