Case Report. © 2013.

TimeKeeper Client Performance and Accuracy

TimeKeeper is the only time synchronization client for Linux that has been rigorously tested for accuracy.

When the network and reference time source are optimal, TimeKeeper can lock to within a few tens of nanoseconds[i]. Network induced delay variations and asymmetries reduce sync quality, but TimeKeeper can compensate in most situations. For example, one FSMLabs customer installation gets sub-microsecond sync between Korea and Japan.

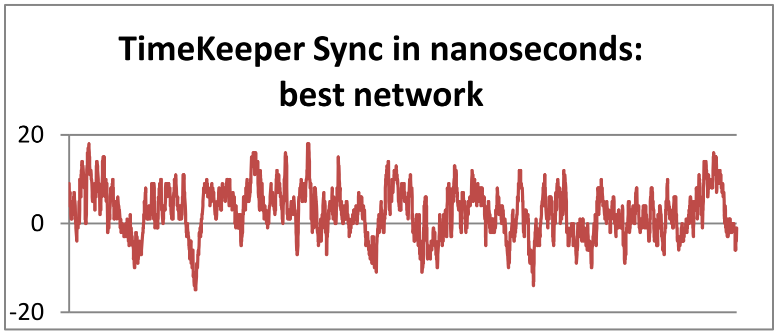

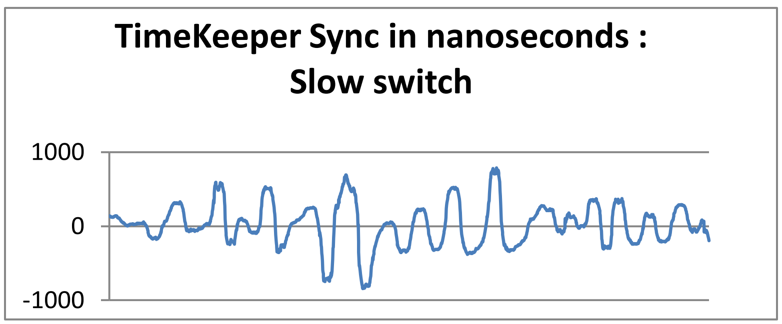

A standard corporate network with complex switch/router structure can easily sync within a few microseconds.

Measured performance of TimeKeeper	
Tested Configuration	**Difference of client time from Reference. Standard Deviation**
Direct (no switch delays)	5 nanoseconds
10GBps loaded network/medium reference sync	280 nanoseconds
Korea to Japan sync	1 microsecond

Even when a non-optimal switch or router intervenes, TimeKeeper can keep standard deviation[ii] within 300 nanoseconds and keep the worst case below 1 microsecond.

TimeKeeper also reduces the cost of accessing time, using kernel bypass technology that extends Linux VDSO. Calls to fetch time are below 10 nanoseconds in most modern server computers.

Test Methodology: External Reference against Internal Sync is essential

TimeKeeper has been tested extensively against "external references" so that its accuracy can be precisely characterized and its self-reporting of sync accuracy can be validated. That is, the synchronized time inside the server is tested against some external reference time , not just against the program's idea of how well it is doing. Without such external test validation, synchronizer client estimates of accuracy are meaningless.

If a client is not properly tracking reference time it won't be able to properly estimate how well or poorly it is doing.

FSMLabs Enterprise Real-Time® Case Report. © 2013.

The free software time clients produce misleading self-reports.

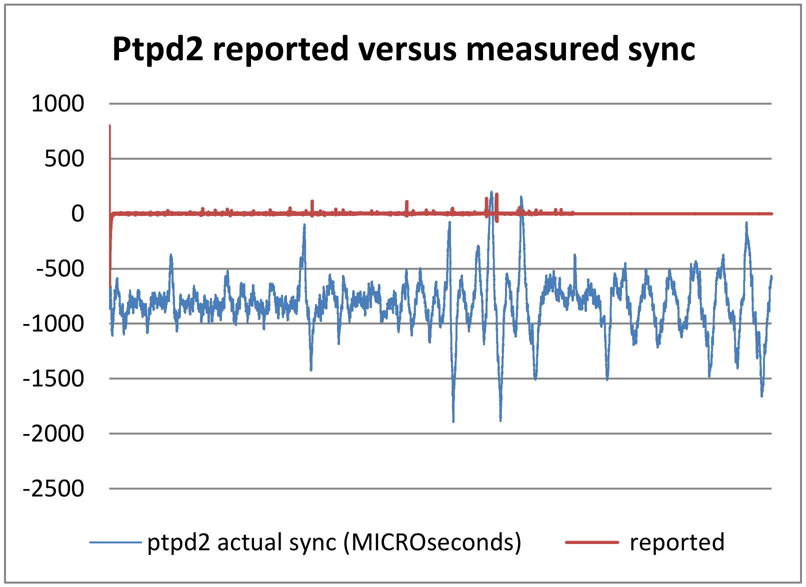

For example, the chart to the left compares reported sync from PTPD2 with measured sync. The measured sync shows a persistent offset of hundreds of microseconds and "jitter" in hundreds of microseconds although PTPD2 reports a one microsecond sync.

The free software clients are capable of producing good sync, but not consistently or reliably and their self-reported performance is not reliable. Those clients will even report high quality sync in the presence of significant failure of the protocol or network adapter hardware.

Internal time: Time at the application

For financial trading, the quality of time synchronization is most important at the application. Synchronizing time at the network port is a simpler problem, but for high frequency trading systems and any high frequency transaction system, the network port time may not correspond at all to application time and having a tightly synchronized network port but a poorly synchronized order routing application is not all that useful. That's why our measurements are of the time provided to the application.

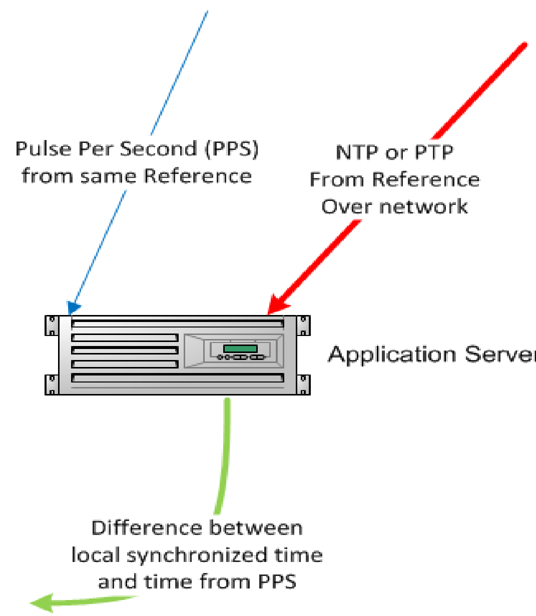

Comparing external time to synchronized time.

TimeKeeper was designed with validation in mind and its unique ability to track multiple time sources independently allows for straightforward testing

A usual test configuration has TimeKeeper synchronize the local clock to some source and also run a secondary synchronized clock using an independent reference. The secondary clock is automatically tested against the synchronized local clock, providing a detailed record of how close the synchronized time tracks the secondary source.

A simple test configuration brings a pulse per second signal (PPS) back from the device that is sending PTP or NTP packets and uses that as the second reference.

The advantage here is that the PPS is being produced by the same device that produces the PTP/NTP feed[iii] so that, at least for some network time devices, both sources are driven by the same clock.

More elaborate tests use GPS bus cards within the application server for the external reference. In this case TimeKeeper reads a time from a bus card register or otherwise accesses the bus card time and compares it to the synchronized time. FSMLabs also uses some tests involve comparing multiple network streams of PTP and NTP against multiple secondary GPS devices.

The following graph is from a test of the second type.

We're displaying the difference between the client local time as generated by TimeKeeper from a PTP link to a network clock and the time read from a GPS bus clock on the client.

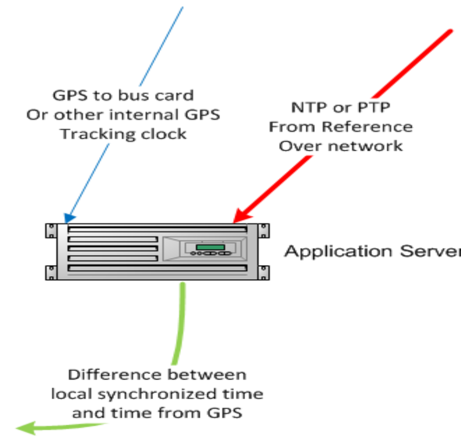

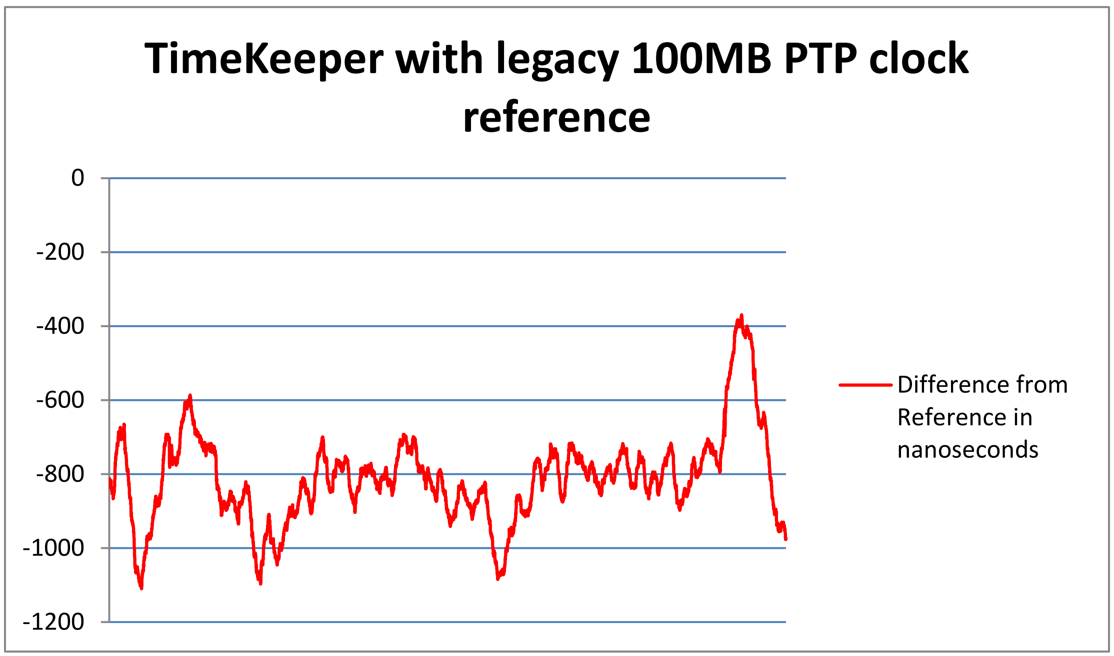

The Y scale is seconds, so we see the worst case difference is close to 1.2 microseconds. This number is actually due to asymmetry in the network caused by the PTP source, which is a legacy network clock with a 100megabit network adapter, plugged into a 1 gigabit switch[iv].

TimeKeeper is compensating but the offset from true time is around 800 nanoseconds.

Time Maps

TimeKeeper includes many capabilities for validating time distribution and finding problems in a time distribution network. The Time Map tab of the TimeKeeper management screen is a data visualization tool which provides a view of the time distributions network in which a synchronization client or a synchronization server (running TimeKeeper server) has been embedded. The length of arcs corresponds to the round trip delay between components. Red arcs carry NTP. Green arcs carry PTP and blue arcs are reference sources (such as GPS). Often we find that the reference time sources are neither as redundant nor as direct as one might want. This kind of map is particularly useful in large corporate networks that are in the process of improving a legacy low accuracy NTP network or mixing PTP and NTP.

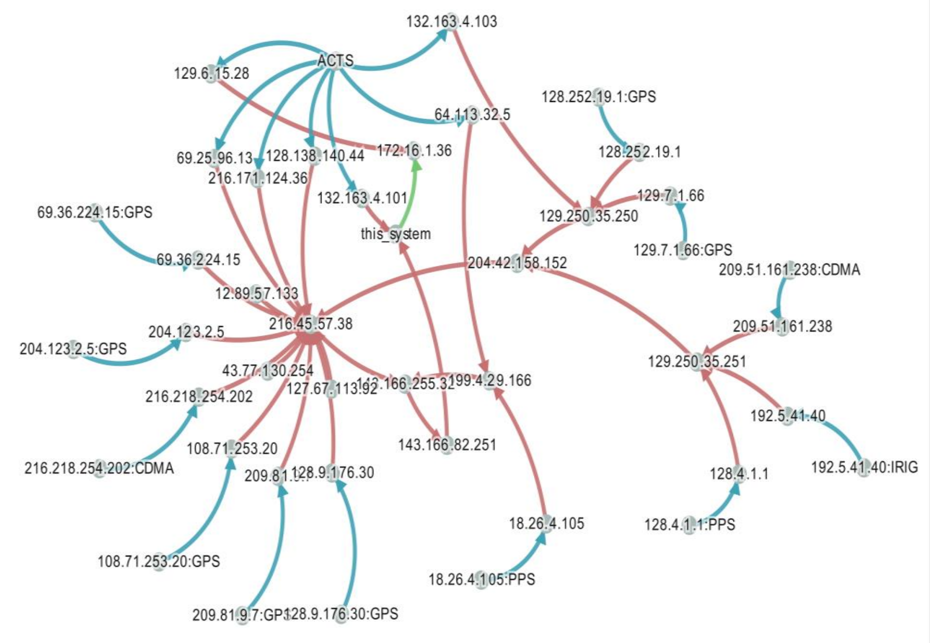

The time map to the right comes from a small corporate network with more PTP connections (IP numbers erased to protect confidentiality). The map shows a PTP server getting PPS (pulse per second) time from GPS and serving PTP to a number of local machines (the cluster in the upper left). It is also connecting to a "bridge" server that has a fallback NTP connection from public NTP server and does a long-haul PTP unicast) over a thousand miles or so to a second cluster which has its own PTP and NTP sources.

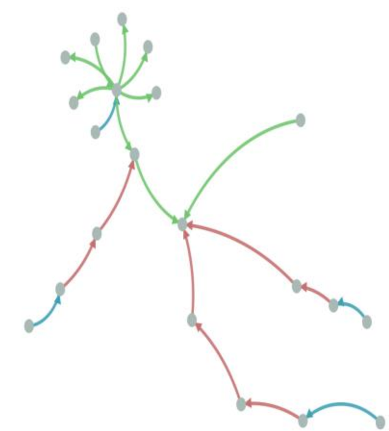

TimeKeeper is particularly useful when you need to glue time distribution networks together - something not unusual for firms engaged in financial trading world wide.

FSMLabs Enterprise Real-Time® Case Report. © 2013.

Source Check measurements

TimeKeeper will also measure frequency information from multiple reference sources and compare them in order to provide some feedback on source quality, to enable automatic fail-over, and to detect errors or compromise (e.g. spoofing) of reference sources. The graph below shows the frequency estimate for the local oscillator from two different reference sources. In this case the differences are within a few picoseconds and validate both sources as reliable.

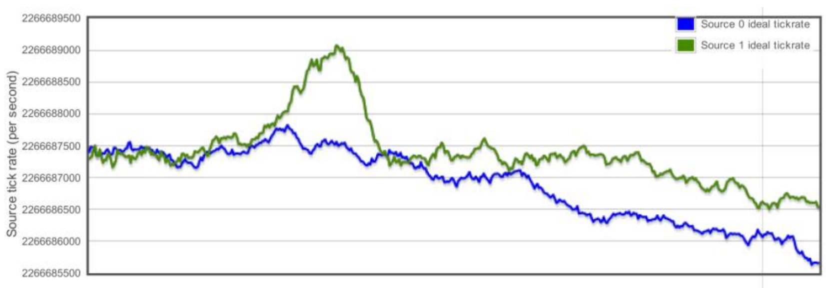

Wrapping it up

FSMLabs can help you structure a test for your own application requirements.

For more information please email sales@fsmlabs.com

Footnotes

---

[i] This test uses TimeKeeper as a PTP source on one network port and has a secondary clock use that PTP feed through a second network port. The secondary clock is driven by the PTP feed and we measure how closely it tracks the primary clock – on the same computer.

[ii] The industry practice is to give time sync quality in terms of standard deviation from the reference time. This is generally a good indication, but sometimes it's also important to know worst case. And it's important to learn the duration of the test. Too short a time can hide weaknesses. Too long a time can bury errors in standard deviation.

[iii] TimeKeeper supports both IEEE 1588 Precision Time Protocol (PTP) or the older Network Time Procotol (NTP). While it is widely claimed that one of these protocols is inherently better than the other, there is no technical basis to that claim.

[iv] [iv] The asymmetry caused by slow network clocks is one reason why FSMLabs is producing 10Gbps Grandmaster Appliances.