FSMTime
SECURE ENTERPRISE CLOCK SYNC™
by FSMLabs

SECURE ENTERPRISE TIME SYNC WHITEPAPER SERIES

Time Protection in Depth at Every Level of the Clock Chain Synchronization with TimeKeeper

Introduction

The TimeKeeper® suite of hardware and software products create a time network that offers highly accurate synchronization, with error-checking, validation, protection and mitigation of time, and the ability to manage and audit the time network. In-depth time protection for assured resiliency is the focus of this paper.

TimeKeeper performs time cross-checks and validations against other NTP/PTP time servers (both TimeKeeper and legacy GNSS grandmaster) for additional self-monitoring and self-healing operation, and can detect outages, failures or time-based cyberattacks to time sources. TimeKeeper will stop using a compromised source and will use the next-best source that it has, which includes a holdover clock state, while the failure is analyzed.

TimeKeeper uses many techniques that are not available elsewhere to validate time at every stage of the clock chain sync – from the top level grandmaster time server to boundary clocks, clients, monitoring nodes, and overall compliance. TimeKeeper detects bad leap seconds, oscillator failures, outages and more subtle problems that include GPS/GNSS jamming and spoofing, and then corrects them based on user configuration. Most of our customers correct these problems by switching to alternative sources or using internal time sources for holdover backup. This allows systems to remain up-and-running while IT personnel are alerted through various mechanisms.

Clock Chain Synchronization

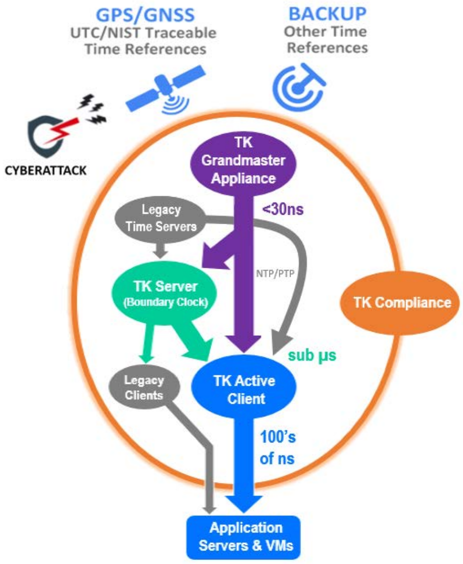

Where bad time comes from and how to mitigate it

Accuracy is built into TimeKeeper from the ground up, and this includes the detection and correction of time ping errors and exploits, including cyberattacks. Multiple patented techniques are used so that threats to keep good time are not overlooked. Some of the errors TimeKeeper protects you from are listed below.

Time protection against network errors

As critical time synchronization increasingly moves to packet based networks with protocols like NTP and PTP, it often shares bandwidth with other traffic. That other traffic can add a heavy load and variations that affect the latency of time sync packets, which directly affects sync quality. Even worse, time sync is typically an after-thought in network design rather than something that's planned for like other network services.

TimeKeeper's intelligent enterprise time sync platform allows your network engineers to view the topology of your clock chain system, by using the TimeMap tool, and determine if any elements have an overly long delay, jitter, or other factors that affect time quality. This allows those engineers to know where time is occurring so that they can resolve those network problems.

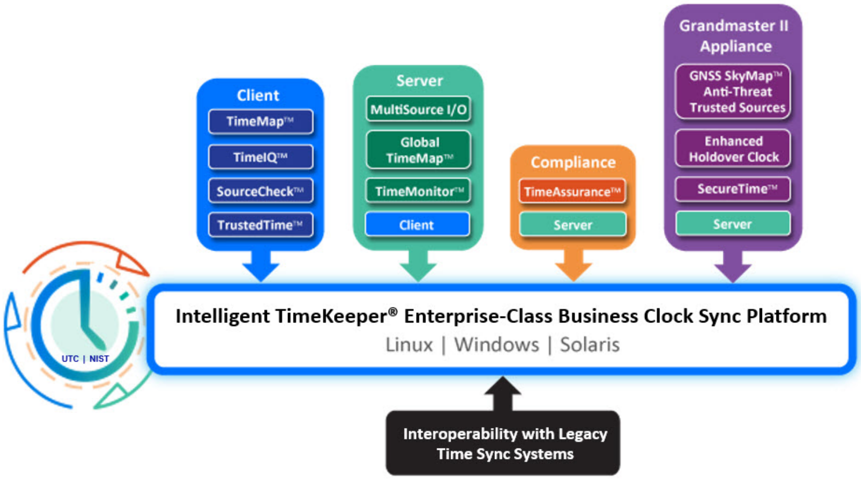

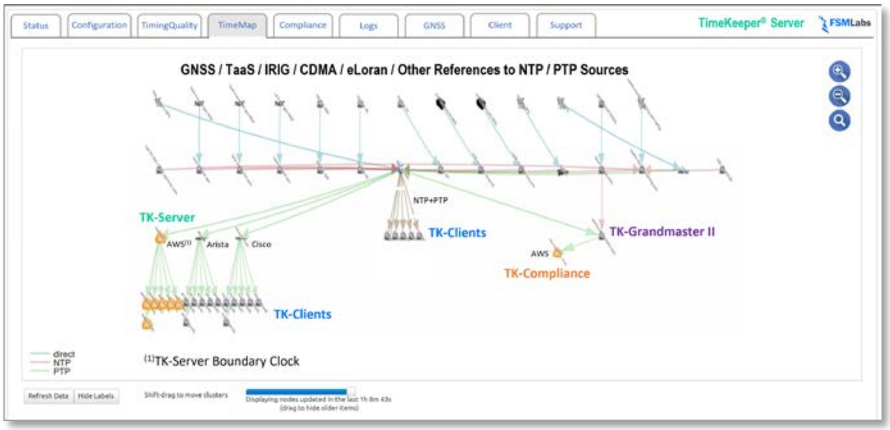

One of the most common problems is when connecting a legacy and slow time server (100 Mbps or 1Gbps) to a 10 Gbps network. Most 10 Gbps switches introduce an asymmetry (or serialization issue) in this configuration that will affect timing accuracy. TimeKeeper Grandmasters provide 10 Gbps interfaces, or up to 100 Gbps if needed, to correct this issue so that interface speeds are matched. Even with legacy time servers that do not have this feature, TimeKeeper provides a means of calibrating the asymmetry and configuring a correction to make up for that error.

Another cause of time errors in networks is from changes in one-way time that if uncorrected can cause jumps of time up to 100s of milliseconds. Asymmetric networks with individual one-way transmit times changing from moment to moment are not unusual. This issue is especially true when connecting data centers through 3rd party network providers. TimeKeeper provides tools for monitoring the one-way transit times of packets over a network and can alert when it senses a change. This capability allows detection of potential time errors and can also alert about other network activity that is being impacted (for example, the latency of market data feeds from a given site increasing significantly). Network engineers typically don't even see this change in latency since other tools that measure, monitor and quantify this issue may not be available to them because of their high cost. TimeKeeper can alert and switch to a good secondary time source or correct for asymmetry to solve these cases.

Time protection against GPS/GNSS obstructions & interferences

Most timing systems get time from a GPS/GNSS source, at the top level of the clock chain synchronization, from a time server appliance, like a Grandmaster or Stratum 1 clock. The TimeKeeper time server appliance monitors GPS/GNSS signal strength and quality to detect multiple types of problems, from benign ones, such as a weak signal due to obstructions, antenna failure or antenna installation issues, to malicious cyberattacks based on GPS/GNSS spoofing and jamming.

TimeKeeper monitors the strength of GPS/GNSS satellite signals in the sky, through the SkyMap tool, to detect obstructions and interference. This is a normal part of any antenna installation since buildings, ACs and other objects can block or interfere with GPS/GNSS signals. Over a few hours TimeKeeper builds a map of the typical signal strength. This allows users to confirm that the antenna is properly installed and the visibility of GPS/GNSS satellites in all parts of the sky is unobstructed. TimeKeeper also monitors the signal strength, and if that signal, in a part of the sky, becomes lower than expected, TimeKeeper will alert you. For example, below is view of the sky where an object is blocking reception to the northwest (image to the left). Then an image showing the same setup when an obstruction was placed on both the western and eastern sky (image to the right).

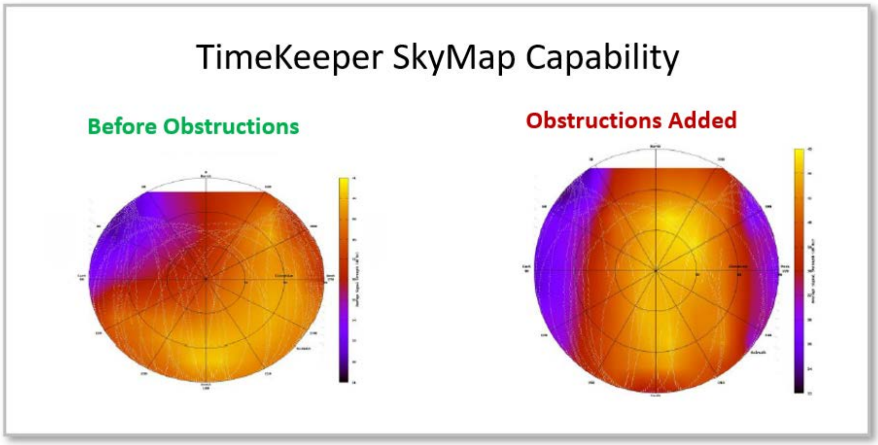

Time protection against GPS/GNSS spoofing & jamming

TimeKeeper can also monitor GPS/GNSS satellite signals for spoofing or jamming. It uses multiple patent-pending techniques in order to do that. Typically jam cyberattacks generate incoherent signals that overwhelm the valid GPS/GNSS signal so that the receiver is unable to read the time. These cyberattacks are well documented and have occurred many times at the London Stock Exchange (see article link) and other locations. One common source of jamming is from (illegal) devices that are available from truck stops and online for $20. The leading-edge TimeKeeper grandmaster time server appliance detects this kind of interference and switches to an internal stabilized oscillator to maintain good time, can be configured to alert you, and will take corrective action if the outage is prolonged, which normally means switching to backup sources.

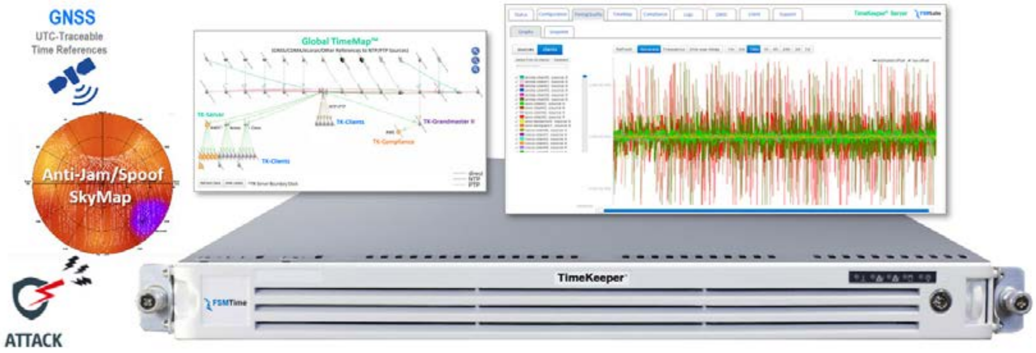

Spoof cyberattacks are more subtle and harder to detect. TimeKeeper has several added patent-pending techniques to handle that. A spoof cyberattack typically works as described here (articles link 1 and link 2). A transmitter is setup on a nearby rooftop, which receives the valid GPS signal from the sky and re-broadcasts that directly to the antenna on a neighboring rooftop that is the target of the cyberattack. Over a period of time, the re-broadcast signal strength is slowly raised until it overwhelms the valid GPS/GNSS signal that the cyberattacked antenna is receiving. At that point the spoofer has full control of the victim's GPS signal since the valid GPS signal is no longer detected by the victim's receiver. Once that's done, the attacker begins to adjust the signal to manipulate time data to accomplish their goals. This type of cyberattack is normally not detectable nor mitigated without using TimeKeeper, an all-in-one integrated grandmaster time server appliance.

To protect from spoofing, TimeKeeper's mitigation technique involves looking at the signal strength in the sky in a similar way to that used for interference detection (described above). In addition to monitoring for changes that represent interference, or lower signal, TimeKeeper monitors areas of the sky for unexpectedly high signal strength. If previously a building, an AC unit or an atmospheric interaction blocked the GPS/GNSS signal, so that sky area or line of sight was weak and it has been that way for weeks, one can expect it to remain the same. However, if that location of the sky begins showing unexpectedly higher signal strength, it can be seen quite clearly within seconds or at worst minutes. That is an indication that a spoofed GPS/GNSS signal is being broadcasted. Similarly, if the GPS/GNSS satellite signal strength appears to be unusually uniform over the entire sky, this can indicate that a spoof cyberattack is underway from a single broadcasted source that is simulating all GPS/GNSS satellites.

It is important to note that this method can detect and mitigate a spoof cyberattack before it has a chance to overwhelm the valid GPS/GNS signal and manipulate the time, or even the position, of your timing system. TimeKeeper can immediately alert IT personnel and then take corrective action.

Time cross-check against false time sources

Another mechanism that TimeKeeper uses to detect and mitigate false time sources (also patent pending) is to compare time among multiple references, including the GPS/GNSS signals, as well as the frequency of those sources (a second level check).That allows detection of a cyberattack to falsify time.

TimeKeeper cross-checks every time source against others (if configured to do so). An over-simplified description of this would be a voting-system for time. If a majority of time sources agree with a time source, then the source is considered valid.

Frequency is defined as the rate at which time progresses; it should progress at the rate of 1 second per second.

TimeKeeper also tests sources for stable and consistent frequency, and a given time source needs to agree with the majority before that time source is accepted as valid. This frequency check can be applied to all time sources – NTP, PTP, GPS/GNSS and anything that TimeKeeper can receive time from in order to detect bad time sources due to network problems, failed hardware, or misconfiguration. It can also detect and mitigate malicious cyberattacks designed to move the clock off of a valid time by using a low rate intended to avoid detection. TimeKeeper can detect these kinds of cyberattacks by comparing both the time (phase) and frequency behaviors of time in order to validate both.

The diagram to the right shows a common setup where a redundant pair of TimeKeeper grandmaster time server appliances in different data center locations validate time against one another. They can also use external time sources including public NTP time servers. Public NTP time servers tend to not be very high quality, but TimeKeeper is still able to extract enough data from them to validate time and frequency for this purpose. An enhanced "SourceCheck" technique validation, at the bottom of the clock chain sync, can be performed on each TimeKeeper Active Client if configured to do so. That way those clients can add an additional level of time security resiliency.

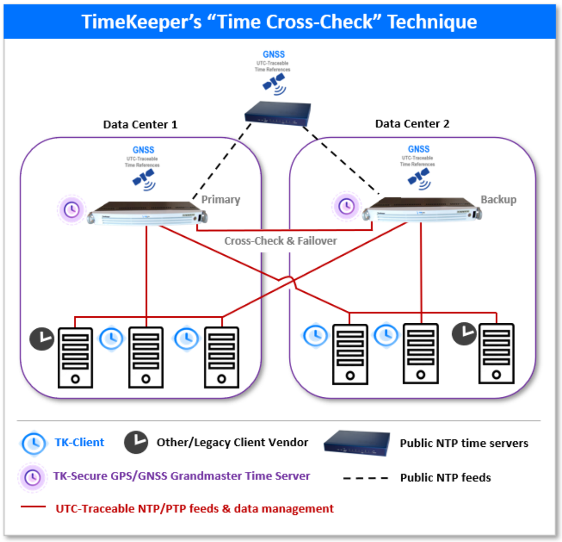
TimeKeeper's "Time Cross-Check" Technique
GNSS
UTC-Traceable Time References
Data Center 1
Data Center 2
GNSS
UTC-Traceable Time References
GNSS
UTC-Traceable Time References
Primary
Backup
Cross-Check & Failover
TK-Client
Other/Legacy Client Vendor
Public NTP time servers
TK-Secure GPS/GNSS Grandmaster Time Server
Public NTP feeds
UTC-Traceable NTP/PTP feeds & data management

TimeKeeper's time cross-check technique like above is critical to avoid a single point of failure bringing down your infrastructure. Deutsche Boerse halting trading for over an hour (Reuters article), because of a time synchronization error that could have been prevented or mitigated with TimeKeeper. This is far from the only example where time caused an exchange to fail.

Time protection against virtual machine inconsistencies

Virtual systems (EC2, Google Cloud, locally hosted virtual and NFV systems) can introduce unusual and unique time inconsistencies that are not common on physical hardware. TimeKeeper is designed to recognize those unique challenges and mitigate them properly.

It's not unusual to see 10 minutes/day error with non-TimeKeeper synchronization technology on virtualized systems. Precise time, given the harsh environment for time sync in the cloud requires the sophisticated algorithms, filters, and network timing model TimeKeeper provides to recognize this environment and handle it appropriately. For example, access to the hardware clock is often not available with these systems, but instead it is emulated in an attempt to maintain a constantly incrementing clock even when a virtual machine is suspended/restarted in order to allow another one to run. So the virtualized clock may perceive that no time has elapsed during one of these events, but in reality many milliseconds may have elapsed. That discontinuity must be mitigated properly in order to both preserve continually advancing clocks that are adjusted smoothly for applications, while still correcting out any errors that are introduced as quickly as possible. This prevents erratic and invalid time on virtual instances that are typically with legacy time synchronization systems.

Time protection against false, missing or extra leap seconds

Many leap second related problems from GPS/GNSS data occurred in 2008 and 2012 that resulted in outages of a number of high profile sites (Reddit, Quantas airlines, and more). TimeKeeper protects you from leap second errors. In 2015, we saw leap second errors in many devices, but because of the effort made by the affected companies' technical staff, most of these problems were mitigated easily.

One error we found was that some units introduced a leap second more than 5 months early (in January), causing many systems to "leap" early and display a 1 second error. TimeKeeper cross-checks and leap second protection safeguards prevented our customers from being affected, because our leap second detection knows the two times of year that a leap second may be introduced and will reject leap seconds outside of that range. Additionally, TimeKeeper cross-checks against other sources verify that the leap seconds that are introduced are valid.

Even if a leap second should occur but is not provided by the time server, TimeKeeper grandmaster will catch that with the time cross-checks technique. The same protections prevent the introduction of double-leap seconds (not uncommon when one time server introduces the leap second and a secondary time server does the same).

Time protection against a grandmaster in "holdover" mode

A high-quality grandmaster time server appliance normally gets time from GPS/GNSS sources. When that fails, due to an antenna being disconnected, GPS/GNSS signals become invalid, so then the grandmaster goes into holdover mode by using a high-precision backup oscillator to keep steady time.

TimeKeeper Grandmaster appliance includes an enhanced double-oven oscillator (DOCXO), and optionally an enhanced Rubidium Atomic Clock (Rb). These oscillators are normally exceptionally stable and hold time accurately for many days, but they do eventually drift. TimeKeeper's enhanced drift rate of the DOCXO is 4µs/day, whereas the Rb's rate is 0.6µs/day.

TimeKeeper software relies on the host system's real time clock, which is generally based on an inexpensive quartz crystal. TimeKeeper software has proprietary technology that regulates the host's clock and gives it a considerably more stable frequency than it would otherwise have.

In the unlikely event of a GPS/GNSS loss that continues for many days, Timekeeper Grandmaster monitors the internal oscillator offset against other (local or remote) NTP, PTP or other time servers. When the internal oscillator time has moved beyond a certain threshold compared to those external time sources, Timekeeper can alert and switch to another time source. This allows the system to take advantage of high-quality oscillators and still catches any unavoidable oscillator drift.

Time protection from protocol failure & spoofing

	TimeKeeper technology is protocol, technology and medium agnostic. This means all the above checks can be applied to any and all time synchronization protocols and sources. This gives you flexibility in setting up your network by bridging disparate time networks (NTP, PTP, PPS and so on), and also allows you to cross check different time sources and formats. The end result is reliability and resiliency in case of time errors or network outages and equipment failures that take out all time networks of a given protocol. For example, it is not uncommon for buggy switches to inject bad time via PTP but not NTP, so TimeKeeper can switch to NTP automatically. The same is true of some PTP enhanced network cards that can inject bad time while letting NTP pass through untouched.
TimeKeeper logging: detect, alert, correct, audit, report	Timekeeper can detect time errors and then take corrective action to switch to properly working time sources. TimeKeeper is also continuously logging time quality data and corrections for all time sources for later analysis and recovery. It alerts when it has detected bad time sources or has taken action, so IT staff can immediately get involved. These logs also provide all data and tools for audits done for time accuracy compliance and reporting, either for internal and/or regulatory use. A single instance of TimeKeeper can also monitor any number of clients and provide metrics, logging and alerting for them all in one place. You can view these metrics and logs on “one pane of glass” or in one set of log files.
Have questions? Simply contact us	TimeKeeper Platform Solution, TimeKeeper Grandmaster II Appliance, TimeKeeper Server or Boundary Clock Software, and TimeKeeper Active Client Software are available from www.FSMTime.com and FSMTime’s resellers. If you have any questions or would like a live demo of TimeKeeper products, simply contact us at sales@fsmtime.com.