

White Paper
Financial Regulations Synchronized with the Latest
Timekeeping Technology



Financial Regulations Synchronized with Latest Timekeeping Technology

The EU's European Securities and Markets Authority (ESMA) and the U.S. Securities Exchange Commission (SEC) recently set new rules concerning the accuracy of time stamps. As markets and market participants are highly interconnected, these rules have a ripple effect on market participants worldwide. Due to the increased need of a more regulated timekeeping system, the EU's Markets in Financial Instruments Directive (MiFID II) has mandated a set of standards to comply to a level of accuracy in all business clocks. This white paper describes a three-faceted timekeeping solution architecture—consisting of a timing master, timing slave, and time monitoring/auditing server—designed to address the growing need for a required level of accuracy. The objective of the three-faceted timekeeping architecture is to map the new financial regulations to the required three technical elements that will allow European trading venues, their members, and participants to establish a network timing solution that is more accurate, reliable, and adheres to the new timekeeping regulations.

Are You Ready for Change?

Time stamps play a key role in reconstructing the chain of events from order origination to execution, and accurate and reliable timekeeping is vital to a fair market. The key time stamping clocks distributed across the network must all reasonably agree if regulators are to audit and correlate market activity at any given point in time. This is particularly important in an era of high-frequency trading (HFT) when markets handle millions of orders, prices fluctuate in microseconds, market participants are widely dispersed, and there are multiple systems in the chain from one end of a transaction to the other.

Meeting the new requirements will force trading venues, investment firms, and other market participants to re-evaluate their timekeeping systems. They will answer the following fundamental questions:

- Do we have the systems in place to assure the required timekeeping accuracy?
- Can we demonstrate compliance?

In many cases, the answer to both questions is probably “no.” The current trading environment is radically different from just a few years ago, compelling firms to adapt to far more exacting timekeeping requirements than ever before. Not only must timing be much more accurate, but that accuracy must be maintained within the high-intensity, widely dispersed market that exists today. As much as the environment has changed, the technology that keeps the time has changed even more. Thanks to new technology and faster systems, trading systems that once were very difficult to synchronize to the levels now required can be synchronized not only with the required accuracy, but also easily and efficiently. System operators will need that efficiency if they expect to comply with the new regulations in time to meet the impending deadlines.

European Regulators Set New Time Stamping Benchmark

In Europe, Article 50 of the EU's MiFID II required that “ESMA shall develop draft regulatory technical standards to specify the level of accuracy” of business clocks. ESMA met that goal on July 6, 2016 when it published its Regulatory Technical and Implementing Standards—Annex I. Within that document (RTS 25), ESMA regulated a technical standard on clock synchronization, mandating Coordinated Universal Time (UTC) as the official time reference for trading systems.

Taking effect as of January 3, 2018, MiFID II's RTS 25 mandates:

- Trading venue operators and their members or participants shall establish a system of traceability to UTC.
- They shall be able to demonstrate traceability to UTC by documenting the system design, function, and specifications.
- They shall be able to identify the exact point at which a time stamp is applied and demonstrate that the point within the system where the time stamp is applied remains consistent.

- Reference time may be UTC disseminated by a satellite system, if any offset from UTC is accounted for and removed from the time stamp.

MiFID II's RTS 25 also mandates the required timing accuracy and time stamp granularity. Automated, or non-manual, systems fall into one of two categories—HFT environments or non-HFT environments, with required levels of accuracy and granularity, as shown in the following table.

Table 1 • Levels of Accuracy and Granularity Required for Trading Venues

Environment	Gateway-to-Gateway Latency Time	Max Divergence from UTC	Time Stamp Granularity
Non-HFT	>1 millisecond	1 millisecond	1 millisecond or better
HFT	≤1 millisecond	100 microseconds	1 microsecond or better

Note: The RTS defines gateway-to-gateway latency as “the time measured from the moment a message is received by an outer gateway of the trading venue’s system, sent through the order submission protocol, processed by the matching engine, and then sent back until an acknowledgement is sent from the gateway.”

The 100-microsecond requirement will be a challenge for some market participants. This is because some participants, such as HFT organizations and exchanges, already synchronize business clocks to more than 100 microseconds. For this reason, the U.S. Securities and Exchange Commission (SEC) is considering amending its plans, on the recommendation of market participants, to also require 100-microsecond synchronization of business clocks to UTC (NIST). It is almost laughable to note that the current SEC synchronization requirement is 50 milliseconds to UTC (NIST) with 1-millisecond time stamp granularity, which is virtually effortless to implement, and correspondingly useless when it comes to effectively reconstructing the audit trail millions of trades happening daily.

Note: UTC (NIST) means the UTC time is traceable to the clocks maintained by the US National Institute for Standards and Technology.

Required Timing Elements: Accuracy, Reliability, and Monitoring

The regulations require trading systems to time stamp key trading event records all within an assured level of accuracy and, in MiFID’s case, an assured level of granularity. The following three technical elements must be present:

- An accurate and reliable source of UTC time.
- Accurate and reliable delivery of UTC time from the source to the time stamping location.
- A way to monitor time accuracy to demonstrate compliance.

System operators generally have two options when it comes to a time source—publicly available time servers or GNSS satellites. Time transfer from Internet time sources can be unreliable and not as accurate because satellite-based time servers reside inside the network. Time from satellite-referenced time servers is a well-proven technology, but proper server configuration is key to maintaining the accuracy and reliability of the time.

Operators also have two options when deciding how to distribute time from the source to trading systems over an internal network—the Network Time Protocol (NTP) or Precise Time Protocol (PTP). Only properly engineered PTP networks referencing time from satellites are well suited to deliver time better than 100 microseconds of UTC. For environments requiring only 0.5- to 2-millisecond caliber timing, operators generally can deliver time from satellites through NTP. This depends on the internal network between the satellite referenced time server and the trading systems.

Besides the choice of protocol, the choice of trading system operating system also affects the achievable level of accuracy, as shown in the following table. Because some operating systems are incapable of keeping accurate and reliable time within the requirements of RTS 25, this reduces the choice of protocols and slaves to PTP (few select Windows versions or Linux).

Table 2 • Time Sync Accuracy when Employing Various Operating Systems and Timekeeping Protocols

Operating System	Protocol	Accuracy Relative to the Source	Excursions
Server 2003, XP, Win8.x, Server 2012, Win10, Server 2016	NTP	>500 μ s	Occasional
Vista, Server 2008, Win7	PTPv2	>100 μ s	Regular
Server 2003, XP	PTPv2	>50 μ s	Occasional
Win8.x, Server 2012, Win10, Server 2016	PTPv2	2 μ s	Rare
Linux PTPd	PTPv2	<100 μ s	Occasional

Note: An "excursion" is when the error in accuracy significantly exceeds preset thresholds.

Because the MiFID II RTS accuracy requirement is a regulation put in place for market surveillance and the desire for a consolidated audit trail, there must be a mechanism to monitor the time on the network and the time stamping engine to demonstrate compliance. This can be challenging and brings to mind the Segal's law adage: "A man with one watch knows what time it is. A man with two watches is never sure." In this case, there will be many clocks that ultimately all need to be tightly synchronized to a single concept of time called UTC. The monitoring system must be able to routinely check the time on the key systems, collect time offset measurements, monitor if clocks are outside the allowed range, alarm on exceptions, and archive all the data in the event of an audit (where proof of compliance is needed).

Distributing Accurate Time can be a Challenge

When regulators stipulate a specific level of accuracy, they are mandating a maximum allowed difference, or offset, between UTC and the time stamp issued by the local trading system. Keep in mind that UTC is an agreement between the world's leading national time metrology labs as to exactly what the time was in the past. There is no singular clock that is the UTC clock. Generally, the atomic clocks that provide the time to the satellite systems, such as GPS, are within a few nanoseconds to UTC and stable enough to assume they will continue to be at any moment in time.

Time servers with satellite timing receivers that synchronize their local clocks from the satellites have evolved over the past couple of decades to become quite accurate to UTC. In fact, the better ones are less than 20 nanoseconds to UTC. The challenge then, as in many cases, is the proverbial "last mile" from the receiver over the network to the time stamping engine. In this last stretch, there are several factors that can add to the time offset between UTC and the time stamping engine clock.

UTC time offsets at the time stamping engine can occur for any of the following reasons.

- Satellite signals are lost and the timing receiver begins to drift.
- Timing packets are randomly delayed in transit across the network.
- The time stamping engine clock drifts between time updates.

These three sources of time offsets will contribute to the offset that authorities regulate (that is, between UTC and the time stamps issued by trading systems). The role that time source technology and NTP/PTP protocols play in timekeeping is to enable correction of all these offsets by providing a time reference source that downstream systems use to steer their clocks.

Here's the challenge: despite these offsets, the time stamped NTP or PTP packets arriving at a trading system must somehow be as closely synced to UTC as if the trading system itself were attached directly to an atomic clock at a national time metrology lab. That challenge is made harder by the fact that packet delays across the local network are not constant. This packet delay variation (PDV or packet arrival time jitter) means that any offset correction technology must not only correct for offsets but also correct dynamically changing offsets in real time.

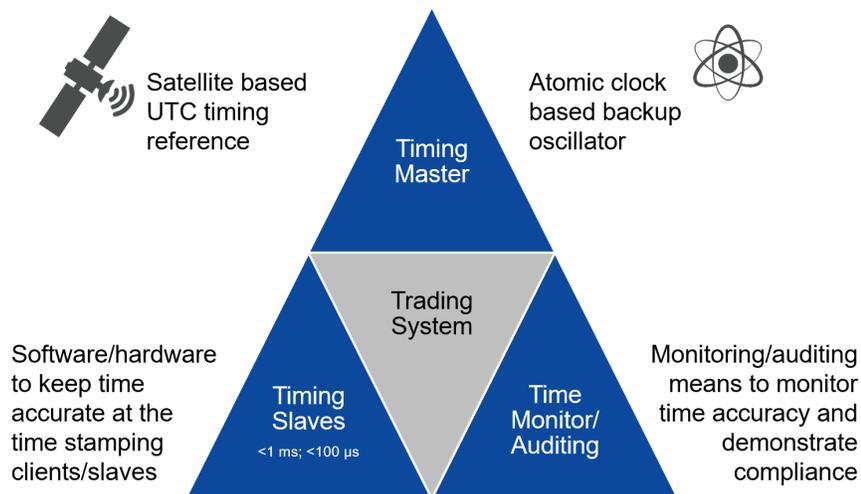
The Compliance Solution: Three-Faceted Timekeeping Architecture

Corresponding to the three technical elements the financial regulations require; the solution is a technical architecture consisting of the following key components.

- A timing master that takes time from the satellite receiver and provides it over the network as PTP or NTP time stamped packets.
- A timing slave—a Windows or Linux computer (typically the trading system itself)—that receives upstream time stamped packets, steers its local clock, and uses that time to time stamp events.
- A time monitoring/auditing server to track timing performance across the network, send alerts when set thresholds are exceeded, and provide an audit trail of timekeeping performance to prove compliance.

The following illustration shows a three-faceted timekeeping solution architecture map to the three technical elements the new financial regulations mandate be present.

Figure 1 • Three-Faceted Timekeeping Solution Architecture



The 100-microsecond to UTC requirement will compel system operators to use a GNSS time reference operating as a PTP grandmaster as the most reliable means to stay well below the required accuracy. This is important if the satellite signal is lost and the overall system starts to drift in time. Use PTP to allow the time stamping system to stay tightly coupled to the grandmaster in the event it drifts. This introduces the requirement for an upgraded oscillator, such as an atomic clock in the grandmaster, to minimize drift and provide the system operator time to address the issue, all the while staying below the 100-microsecond to UTC requirement.

With PTP, the PTP slave might use either hardware or software time stamping. Hardware time stamping eliminates asymmetric delays attributed to timing packets transiting the operating system stack of the slave computer responsible for the time stamping. Using a PTP-capable network interface card (NIC) is the typical deployment scenario for single-digit microsecond accuracy or better time stamping activities. For time stamping in the 10's of microsecond accuracy range, users can run either Linux PTPd open source software or a Windows-based PTP client.

Lastly, a PTP-enabled monitoring/auditing solution will monitor and audit time offsets, and issue operator alerts when large offsets occur. The key here is a system that is synchronized to UTC, checks the time on the time servers, and collects the time correction information from the slaves. By being synchronized to UTC, it can cross-check the time on the GPS-referenced time servers against other time servers and its own internal clock. All offsets are logged for later retrieval and large offsets can trigger alarms. The audit server is not required to have a better time than the GPS-referenced time servers, but rather close in time and simultaneously checking against multiple time sources for general correctness and compliance. The audit server must also routinely and periodically collect all time offset/correction data from the slaves to archive the data, alarm on exceptions, and provide reporting capabilities.

Arriving at the Best Solution

System operators will need to decide which specific hardware and software features to employ to implement timing masters, timing slaves, and the time monitoring/audit solution. For timing masters, those decisions affect the time stamping speed and accuracy even if satellite reception is lost.

Timing Master Decisions

- **Satellite systems supported:** The more satellite systems supported (for example, GPS, GLONASS, BeiDou, and SBAS), the more likely satellite reception will be maintained if reception from one of the systems becomes unavailable. Multi-constellation support also aides in improved time accuracy in limited sky-view environments such as urban canyons.
- **Type of internal clock:** A rubidium atomic clock will drift less than one microsecond in one day and less than three microseconds in three days, assuring a highly accurate time source if satellite reception is lost. This provides the operator time to address the issue, while remaining compliant to the 100-microsecond to UTC requirement.
- **PTP packet time stamping capacity:** Depending on the environment size, demand for PTP packets may be very high. The time server must be able to accommodate the load and support the appropriate PTP profiles best-suited to the network topology.

Timing Slave Decisions

- **Level of sync accuracy:** Operators will also want to assure the trading system (PTP slave) is on an operating system capable of keeping the time to better than 100 microsecond (as shown in [Table 2 \(see page 3\)](#)). The PTP slave's role is to steer its host clock to keep synchronized with the master, it should do so with high accuracy. Consideration will need to be given to PTP hardware time stamping NICs or merely just PTP software based solutions. The popular choice on the Linux side is the open source PTPd. For Windows, operators should look for a third-party proprietary solution that delivers high-accuracy synchronization, something the Windows Time Service (W32Time) was not designed to provide.
- **Flexible network topology:** Consideration should be given to the capabilities of the switches/routers between the master and slaves and if they support PTP. Switches with PTP boundary clock or transparent clock functionality can be useful in reducing the time accuracy degrading effects of asymmetric path delays caused by queuing.
- **System impact:** Ideally, the solution would run in the background, consuming very little resources, and not interfere with any critical system executables.

Monitoring/Auditing Decisions

For monitoring and auditing, operators need to consider features that reduce the compliance burden and enhance reliability.

- **Centralized control or just monitoring and auditing:** System operators may want to leverage the solution as a single point of control to accomplish tasks such as installing, upgrading, and configuring PTP slaves.
- **Single record repository:** The ability to see what's happening across the timekeeping infrastructure and record the events, such as when slaves fall out of sync. Then record it centrally for auditing purposes (crucial for compliance).

- Easily tailored alerts and reports: Different environments call for different alert thresholds, report formats, reported data, and so on. This means the ability to accommodate a market participant's individual needs will greatly enhance speed and thoroughness of compliance, while reducing management overhead.
- Auto failover and recovery: This enables sites to stay compliant if one of the monitoring/auditing servers goes down.

Navigating the Path to Compliance

Trading venues, investment firms, and other market participants have critical decisions to make if they expect to comply with the MiFID II requirements and forthcoming SEC mandates by the rapidly approaching deadlines. As outlined previously, a three-faceted timekeeping architecture (consisting of a timing master, timing slave, and time monitoring/auditing server) is key to deploying a network timing solution that will be accurate, reliable, and compliant to the new regulations. The architectural choices you make will have long-term impact on your ability to comply, making it imperative to choose wisely now.

If you're facing the challenge of deploying a MiFID II or SEC-compliant timing infrastructure for your trading platform, Microsemi has the in-house timing expertise to recommend the appropriate architecture to meet the requirements for the master, slave, and monitoring/auditing components. Ultimately, you will want a timing system that is accurate, reliable, secure, and monitored for peace-of-mind trading compliance. Microsemi can help you get there.



Microsemi Corporate Headquarters

One Enterprise, Aliso Viejo,
CA 92656 USA
Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Fax: +1 (949) 215-4996
Email: sales.support@microsemi.com
www.microsemi.com

© 2017 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.

Microsemi Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; enterprise storage and communication solutions; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, California, and has approximately 4,800 employees globally. Learn more at www.microsemi.com.