

MiFID II

CLOCK SYNC YOUR DATA

BEST PRACTICES FOR IMPLEMENTING RTS-25

BY RAYMOND RUSSELL CTO, CORVIL



CONTENTS

- INTRODUCTION 3
- EXPERIENCE WITH UTC TIME SOURCES, DISTRIBUTION AND TIMESTAMPS 4
 - Clock Discipline Processes 5
 - Timestamp Measurement Methods 6
- MEASUREMENT OF TIMESTAMPED TRADE EXECUTION EVENTS 7
 - Timestamp Data Sources 7
 - Understanding Causality Issues 8
- EXAMPLES OF MIFID II DATA EVENTS REQUIRING CLOCK-SYNC 10
- PRACTICAL CONSIDERATIONS FOR IMPLEMENTATION 11
- SUMMARY 12

INTRODUCTION

MiFID II requirements for clock-synchronization, as summarised in RTS-25, require firms and venues to timestamp events accurately relative to Coordinated Universal Time (UTC) and to an appropriate level of granularity depending on where a trade is executed and the gateway-to-gateway latency of the venue. Specifically, Article 4 of RTS-25 states “Operators of trading venues 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, functioning and specifications. They shall be able to identify the exact point at which a timestamp is applied and demonstrate that the point within the system where the timestamp is applied remains consistent. Reviews of the compliance with this Regulation of the traceability system shall be conducted at least once a year.” ESMA/2015/1909 Section 3.1 further clarifies the use of UTC clock-synchronisation for “reportable events”.

This e-book examines the implementation considerations for deployment of robust timekeeping at acceptable locations (application, host and wire) and measurement methodologies (software and hardware) that will meet ESMA’s requirements and forensic intent while giving investment firms greater flexibility, lower complexity and lower costs to fulfil their obligations in a timely manner.



EXPERIENCE WITH UTC TIME SOURCES, DISTRIBUTION AND TIMESTAMPS

Over the past several years we have worked with the leading venues, participants and market makers throughout the world to develop robust, cost effective and accurate methods to record the timing of electronic trading events. We have found that accurate and precise timestamps of electronic trading events are the critical first step in making sense of these events in terms of latency measurement, forensic investigation and surveillance of trading activity.

NTP strengths

- Mature protocol
- Most systems use the same implementation for divergences within 1 millisecond

PTP strengths

- Standardises aspects of network time-synchronisation for much tighter divergences than NTP
- Can achieve well below a microsecond of divergence with right hardware deployment

NTP weakness

- Requires extensive modifications for tighter divergences
- Modifications may be difficult to support

PTP weaknesses

- Newer protocol still ironing out wrinkles in early implementations
- Leaves the disciplining of local clocks up to the implementation
- Implementation of clock disciplining can vary widely in sophistication
- May require parallel network for PTP traffic, HW support from switches

divergences specified in MiFID II (100 microseconds). However, guaranteeing bounded divergences from UTC, as required by RTS-25, is far from easy. NTP and PTP each have strengths and weaknesses when it comes to reliably and consistently achieving the bounded divergences specified. Reliability and consistency of reporting should be a key requirement for any regulatory reporting system.

Given these strengths and weaknesses, we have found that it is often not sufficient to deploy a network-based time-keeping protocol and assume it will work correctly.

A further time-distribution option is PPS (Pulse Per Second) which is an analogue electrical signal usually driven directly from a GPS receiver. PPS offers high reliability and accuracy, but is challenging to distribute widely and can only be consumed by specialised network cards.

Corvil analytics consume timestamped data from our own appliances and from other systems, such as packet brokers. We support PPS and PTP synchronisation, and can verify the stability of PPS and PTP time sources. Our analytics must sequence the timestamped data correctly to deliver latency measurement, forensic investigation and surveillance of trading activity. We have learned through experience that it is critical to constantly run sanity checks on all sources of timestamps, even when all the sources are supposedly well-synchronized to tight divergences. Without those sanity checks it is difficult to determine if odd event sequencing occurrences are the result of normal clock discipline processes or not.

Time-distribution technologies such as NTP and PTP are readily available today. Both technologies can deliver time-signals to well within the tightest

CLOCK DISCIPLINE PROCESSES

When NTP or PTP are well-deployed, local system clocks are routinely checked against the master clock and corrected (disciplined) when the local clock has diverged from the master clock. It is normal for network congestion and other noise to occasionally cause a local system clock to drift far (relatively speaking) from the master. When the local system realises its clock has run too far ahead of the master most implementations of clock-discipline will jump the clock back to being as close as possible to the master.

In some respects, this is a sensible action since it corrects timestamping errors as quickly as possible. However, this can result in a compromised history. An event that occurred before the correction of the clock can end up with a later timestamp than an event after the correction. In the worst case, the reporting could make it look like the t5 event happened before the t4 event, when the opposite is true!

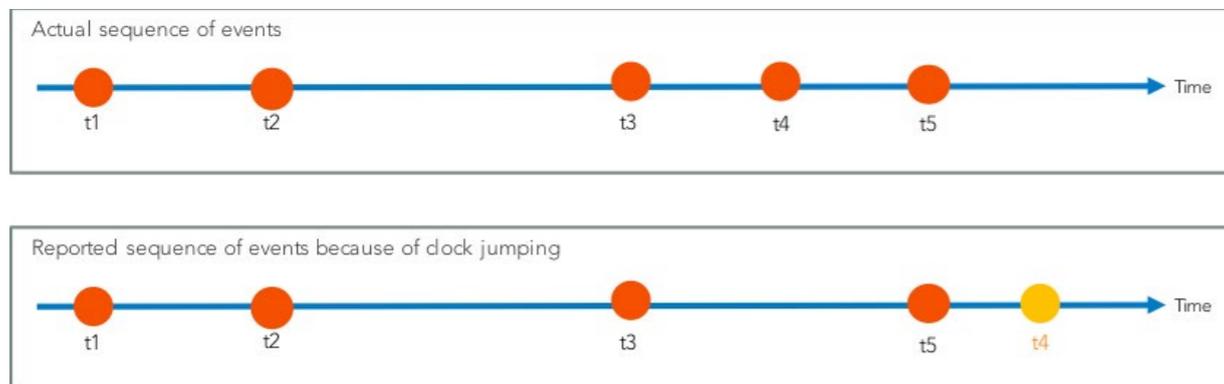


Figure 1: Effect of clock jumping on event reporting.

These jumps in clocks are usually very small, on the order of microseconds. While these jumps are not noticeable on a human scale, our experience has shown that they can skew the reporting of critical events and forensic conclusions. Now that MiFID II is mandating microsecond clock synchronization for some events, such clock jumps must be well understood to properly cleanse event reporting.

The practical lesson here is to double-check your selection and configuration of timestamping technologies. By understanding how your NTP or PTP implementation is disciplining its clocks one can put in place appropriate sanity checks on the timestamps being recorded.

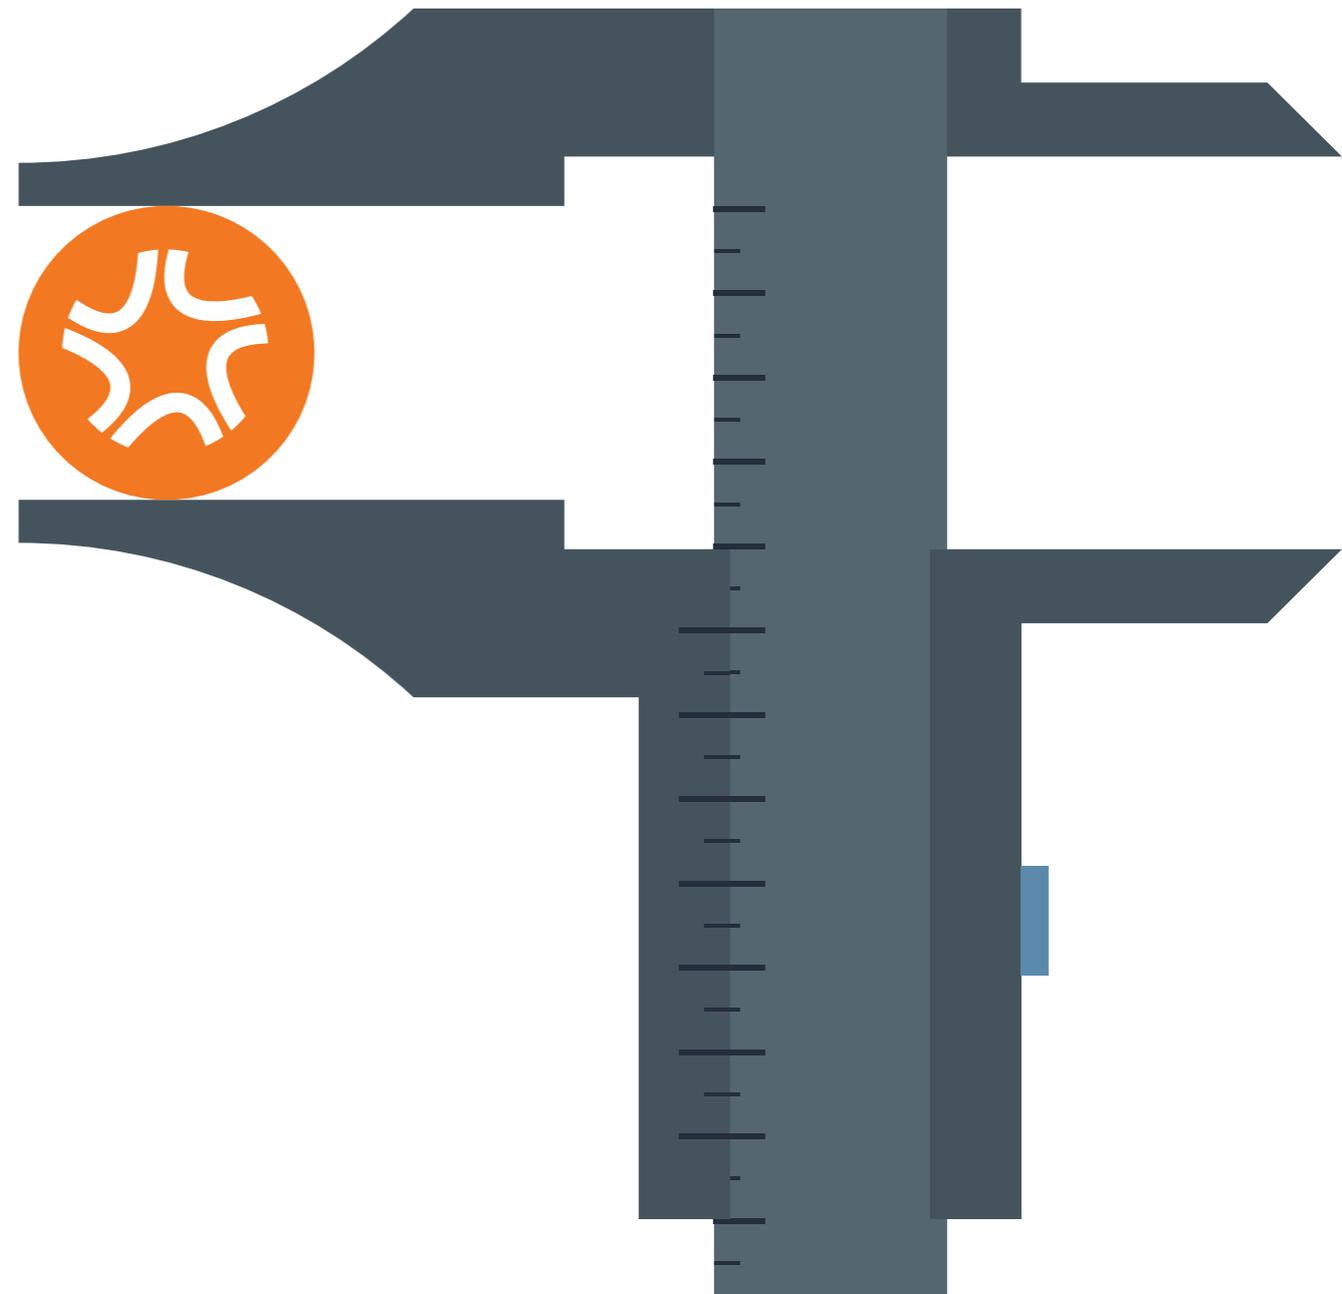


TIMESTAMP MEASUREMENT METHODS

Methods for measuring and applying timestamps to electronic trading events fit into two broad categories:

1. **Software Timestamps** – these are timestamps applied to a specific event of interest using software running on a host machine that is synchronised to a reliable clock source.
2. **Hardware Timestamps** – these are timestamps applied to a specific event of interest using specific purpose hardware in a network switch or a network interface card (NIC) in a host machine.

In general, hardware based timestamps are more reliable, have higher time precision (i.e. nanoseconds) and are less ambiguous. However, some complex trading events involving decision-to-trade and matching-engine events require the use of software timestamps that are accurately synchronised to a UTC time source.



MEASUREMENT OF TIMESTAMPED TRADE EXECUTION EVENTS

TIMESTAMP DATA SOURCES

The two timestamping measurement methods (hardware and software) can be used to create three timestamp data sources for trade execution events:

- **Application Timestamps** - These are timestamps made inline by application software running on a host machine.
- **Host Timestamps** – These timestamps are made inline by system software running on a host machine.
- **Wire Timestamps** – These timestamps are made passively by hardware that takes a copy of the network packets, timestamps the packets and then decodes the underlying messages to recreate the trading context of the traffic, e.g. orders and/or market data.

In practice the three datasources lend themselves to specific use cases.

Application Timestamps are generally used to record the time at which a specific decision was made by the trading application. For example, the time at which a matching engine made a trade. This event time would be when the matching engine software made a match decision between buyer and seller. Typically, this time is recorded in the log file and inserted into the order execution response message for third parties to retrieve and use in their operations.

Host Timestamps are generally used to record the time a specific message was received or sent by the host machine. For example, the time a gateway receives an order or sends an order response. The exact location where this measurement is made in the host stack can vary. It can happen at the socket layer or deeper within the software stack. This can introduce some level of variability and ambiguity in comparing event times between host machines.

Wire Timestamps are most often used to unambiguously determine when a specific message was received and when it was sent by a trading function. For example, wire timestamps are frequently used to unambiguously determine when a gateway received an order or when it sent an order. In practice, wire timestamps tend to be more reliable and precise as there is no variability and ambiguity as to where in the host stack the measurement is made. Reducing measurement variability and ambiguity becomes important when reconstructing the order of events across multiple systems. Wire timestamps are broadly used throughout the industry for troubleshooting potential causality issues with application and host based timestamps.

UNDERSTANDING CAUSALITY ISSUES

The intent of MiFID II timestamping requirements is to forensically verify and survey the specific sequence of reportable events leading to a transaction outcome, so it is important to understand the causality relationships between the available timestamp data sources.

Figure 2 shows a general picture of a host machine processing incoming events from the network wire, processing them and responding with a sequence of outgoing events. This is typical of the situation we encounter with trading gateways, smart order routers and matching engines. In this example, the host machine is UTC synchronised via PTP distribution of the clock signal to the host machine.

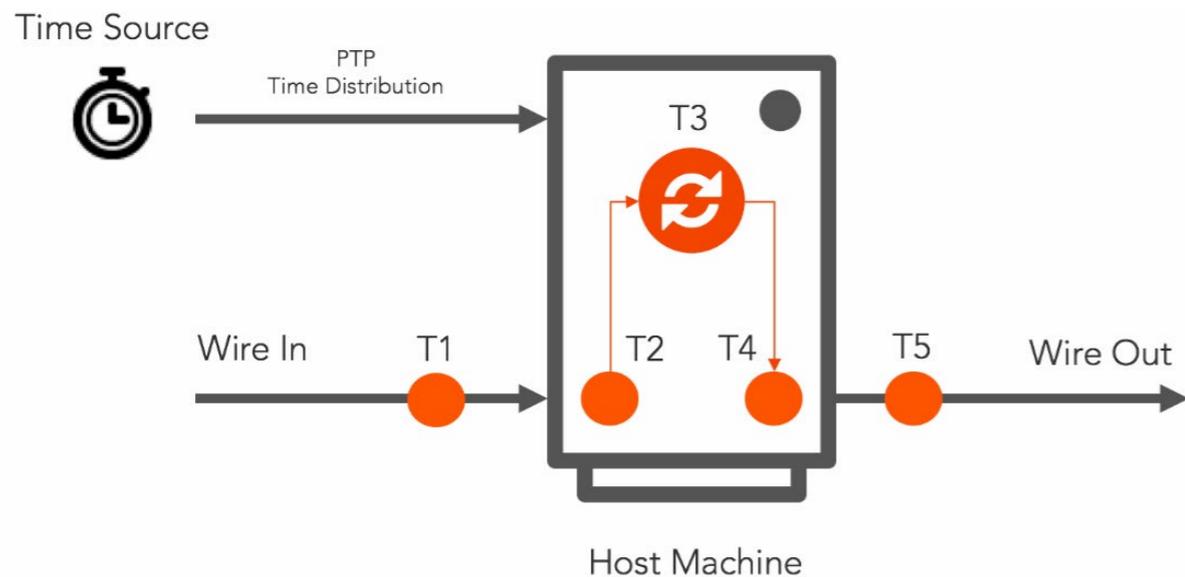


Figure 2: Wire, host timestamps, and application timestamps for measurement and recording of reportable electronic trading events.

The figure also shows the set of timestamp measurements that can be made using the three timestamp sources outlined above. The timestamp measurements are:

- **T1** – a hardware measured timestamp from the wire of the receive time of the specific event of interest e.g. order received, tick received. T1 is the earliest possible time for the event receive time.
- **T2** – a software measured timestamp from the host of the receive time of the specific event of interest e.g. order received, tick received.
- **T3** – a software measured timestamp by the trading application of the decision time for the specific event of interest e.g. decision to trade, decision to match, decision to route.
- **T4** – a software measured timestamp from the host of the transmit time of the specific event of interest e.g. order sent, tick sent.
- **T5** – a hardware measured timestamp from the wire of the transmit time of the specific event of interest e.g. order sent, tick sent. T5 is the latest possible time for the event transmit time.

Because we know that an output response cannot happen before an input event, we can reliably say the following about the above set of event times.

- T2 should always be later than T1
- T4 should always be earlier than T5
- T3 should always be later than T1 and T2
- T3 should always be earlier than T4 and T5
- T1 and T5 represent the lower and upper bound of the event decision time

In other words, the causality relationship can be represented as:

- $T1 < T2 < T3 < T4 < T5$

The reported timestamp data will always match this causality relationship – if all measurements of T1 to T5 are performed accurately and consistently by their respective timestamp data sources.

In practice we find that this is not always the case. Sometimes we find that host and applications timestamp measurements are not implemented properly or there is a problem with the host clock source. For example, where the host clock was a few microseconds ahead of the correct time, the application, host and wire timestamps could report data that looks like:

- $T2 < T3 < T4 < T1 < T5$

In other words, the reported data would imply that the T4 event (for example, outgoing order) was not caused by the T1 event (for example, incoming market data tick), when the opposite is true!

Wire timestamps are typically implemented using specific purpose hardware and tend to be more accurate and reliable. For this reason, they are often used to calibrate and troubleshoot potential problems with host and application timestamp measurements to assure correctness and fidelity.

EXAMPLES OF MIFID II DATA EVENTS REQUIRING CLOCK-SYNC

The following are example events that must be tracked and/or reported to competent authorities under the new RTS requirements:

- **Time of receipt of order and Time of forwarding of order** - RTS6 ([2], p234) requires order record keeping for algorithmic trading environments to record "The exact date and time of the receipt of the order or the exact date and time when the decision to deal was made."
- **Information relating to outgoing and executed orders** - RTS6 ([2], p240) requires order record keeping for algorithmic trading environments to record "The exact date and time of the submission of an order to the trading venue or other investment firm" and "The exact date and time of any message that is transmitted to and received from the trading venue or other investment firm in relation to the order."
- **Gateway to gateway latency** - RTS7 [2, p262-3] requires venues to measure the "time delay between receiving a message in any outer gateway of the trading system and sending a related message from the same gateway after the matching engine has processed the original message;". RTS25 ([2], p504) clarifies this as "Gateway to gateway latency shall be 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".

Consideration of how one can best implement compliance with these new rules will depend largely on the current state of the IT systems supporting the investment firm.



PRACTICAL CONSIDERATIONS FOR IMPLEMENTATION

To comply with the requirements specified in ESMA/2015/1909 [1] and ESMA/2015/1464 [2] and address the regulator's forensic intent, trading venues and participants will need to instrument their IT systems with new technology. The complexity and cost of this undertaking will depend on the approach taken by the investment firm and the current state of its server systems and software architecture.

For example, investment firms that choose to retrofit their existing servers, application stacks and databases with support for microsecond or nanosecond precision clock-sync'd data measurement are likely to incur significant deployment effort, including:

- PTP distribution to every server – this may require a hardware upgrade of every switch to ensure the necessary network-level hardware support required for reliable distribution of the time service.
- Server upgrades to install necessary PTP hardware (PTP-enabled NIC if not supported on motherboard management port).
- Application code changes to dozens or hundreds of applications and databases, to implement host and application level timestamps with the required accuracy and precision.

By contrast, leveraging wire timestamp data sources where applicable are likely to have less complex, more reliable, less costly projects that are faster to implement due to:

- Reduced size of the PTP distribution infrastructure - each wire measurement appliance or timestamping aggregation switch can monitor data to hundreds of servers.

- No change necessary for server systems, trading application software and databases as timestamping and reporting is offloaded to a non-intrusive wire measurement system.
- Aggregation of data from many servers with a single network monitoring point, reducing the complexity of data aggregation, normalisation and sequencing.

For these reasons we are advising investment firms to fully consider where they can use wire timestamp measurements of reportable events to meet the new RTS rules and minimise the need to make changes to their applications and databases.

In addition, we recommend that investment firms future proof their MiFID II implementation by targeting sub 10 microsecond UTC divergence and sub 10 nanosecond timestamp granularity. Currently RTS-25 specifies a granularity of "1 microsecond or better". Given the regulators intention of using this data for forensic evidence of market abuse, the currently specified accuracy and granularity is insufficient to unambiguously determine sequence of events and causality in a reliable fashion across venues and market participants. In addition, the response time of venues will continue to get faster over time, exacerbating the problem. Therefore, it is prudent for investment firms to get ahead of this issue.

SUMMARY

Our experience suggests that investment firms should:

- Carefully select and design UTC timekeeping systems and devise diagnostics to continuously assure system accuracy
- Use independent, non-intrusive wire timestamp measurements of clock-sync data where possible
- Use lightweight agent instrumentation of application stacks to offload overhead and minimise re-design of existing code
- Future proof your implementation by aiming for sub 10 microsecond UTC divergence and sub 10 nanosecond timestamp granularity

This overall approach will minimise cost and complexities while fully complying with the specific rules and intentions of ESMA, today and into the future.



Corvil safeguards business in a digital world. We see a future where all businesses trust digital machines to algorithmically conduct transactions on their behalf. For some businesses, this future is now. We provide big data analytics products that examine digital machine communications, in machine time, and apply analytical and statistical methods to deliver new levels of trusted, streaming intelligence needed by business, IT and security operations teams to safeguard the transparency, performance, and security of critical business applications and services. Corvil was forged on Wall St where it is trusted by leading financial institutions to safeguard their businesses in a digital world.