VIA ELECTRONIC MAIL

January 7, 2014

Marcia E. Asquith
Office of the Corporate Secretary
FINRA
1735 K Street, NW
Washington, DC 20006-1506

Dear Ms. Asquith:

RE: FINRA Regulatory Notice 14-47: Request for comment on a Rule Proposal to Tighten Business Clock Synchronization Requirements.

FINRA published Regulatory Notice 14-47 ("RN 14-47") on November 3, 2014 requesting comment on a proposed rule to implement the Tighten Business Clock Synchronization Requirements. FSMLabs appreciates the opportunity to comment on the proposed regulation and commend FINRA for its attention to this important issue.

In our view:

- The 50 milliseconds requirement for electronic clocks is a major improvement and can be met with low cost off-the-shelf software only.

- A 1 millisecond requirement would not impose significant additional costs on market participants.

- 1 <u>micro</u>second is practical with low-cost off-the-shelf technology.

- Assurance, reliability, and traceability are critical issues. Even the current 1 second standard is often violated and many market participants do not have the ability to validate timing accuracy or to alarm on synchronization failure or even to support meaningful forensics in failure situations. This is despite the existence of low cost, off the shelf technical solutions.

- Long term data integrity will require that the 50 millisecond standard be a first step. Tighter standards will be necessary in the near future.

FSMLabs develops and markets time synchronization technology, software and hardware that is widely used in the financial trading industry, which places us in a knowledgeable position to comment on the feasibility of implementation for this proposed rule. Data for the Security and Exchange Commission (SEC) Midas system is collected on a computer network that uses FSMLabs' technology. There may be alternative methods to address the issues we highlight here, but our technology provides production solutions that are easily affordable and use off-the-shelf technology. To follow are our thoughts on why FINRA's proposal can be easily implemented by financial services organizations as outlined in the proposal and the core issues and risks related to data governance driving the need for regulation to address the issue.

At the end of this letter, we've provided an appendix which includes a glossary and additional background on clock synchronization.

Meeting the 50 millisecond standard and beyond.

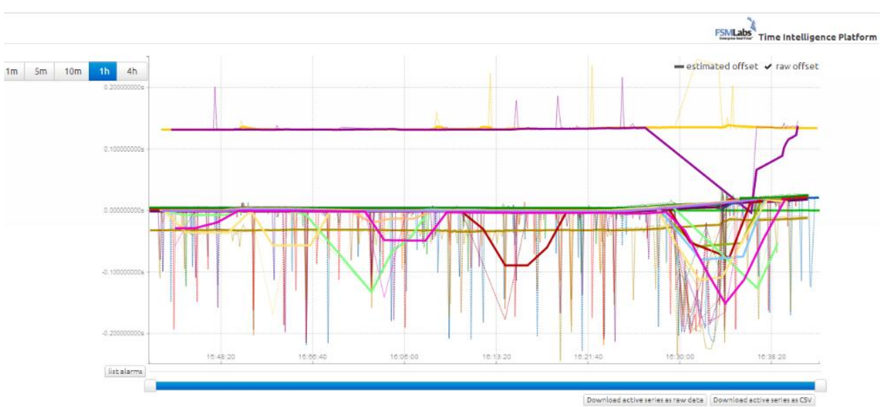

Client Time Synchronization Software can get time over the open Internet from National Institute of Standards and Technology (NIST) servers that is sufficiently accurate to meet the 50 millisecond standard. Multiple sources are available, so sufficiently capable client software can implement cross-check, failover, and timestamp integrity monitoring.

Even cloud based clients can meet the millisecond standard.

Many systems already have internal requirements for better than 50 milliseconds. High-accuracy time is obtained by pulling time from GPS satellites or NIST TMS boxes and distributing that time within data centers and cluster. The critical need for those systems is reliability and assurance.

Assurance, Data Integrity, and Data Governance

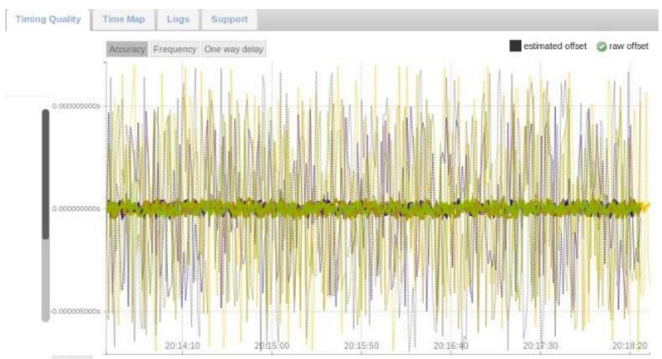

The main weakness of many existing clock synchronization technical systems is in Data Governance, specifically in assurance and traceability. Organizations need to be able to answer a number of questions, such as:

- What processes are in place to monitor time synchronization quality – both technical and operational?
- Does the time synchronization technology in use provide:
    - Cross-check of multiple sources?
    - Automatic failover?
    - A traceable audit record?
    - Alarms and alerts?
    - Reliable self-test (when it says it is within 10 milliseconds, is that reliable)?
    - Simple configuration, management, and provisioning (will it be, in practice, installed and operated properly)?
    - An upgrade path to better than 50 milliseconds?

- Do support staff have either sufficient time synchronization expertise or access to support from a qualified source?

Without good answers to these questions, market participants might, in theory, meet some requirement for synchronization, but may not meet it in practice.

Data Governance and Management

The biggest driver of costs in time synchronization and the greatest source of technology failures is data governance, specifically lack of comprehensive management of goals and risks and proper analysis of costs. Because clock synchronization is often considered a "down in the weeds" technical issue instead of a business logic requirement, it is often shuffled off to engineering staff without being integrated into general technology capability and risk management. Improvised solutions do not achieve accuracy goals and lead either to a sequence of increasingly costly "band-aids" (such as scripts to check for problems) or to wishful thinking (such as choosing to believe ntpq numbers are reliable). Clock synchronization, particularly in large networks is technically complex. The technical papers from IMC on their in-house clock synchronization (Estrela & Bonebakker, 2012) efforts showcase that a specialized, highly skilled IT team and significant technology budget may not be enough to provide robust solutions in the absence of clear requirements.

The call for comments includes this note:

> *In this regard, FINRA notes that the range across market participants could in fact be twice as large as the allowable drift.*

For example, if one firm's clock is 50 milliseconds behind and another firm's clock is 50 milliseconds ahead, the variance between events reported by these firms could be 100 milliseconds. Accordingly, FINRA believes it is important to set the shortest allowable drift that is reasonable and can be achieved by the majority of firms.

This is a key point, and it makes fault tolerance and forensics even more critical. If either party lacks solid cross-check capabilities, even the 100 millisecond variance is not assured.

Fault Tolerance and Traceable Audit

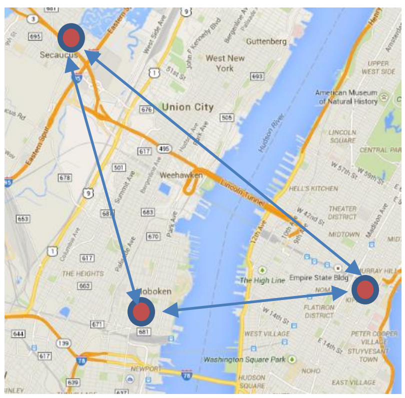

Any basis for data integrity, reliability and traceable audit in clock synchronization must begin with multi-source client capability on application servers. In the absence of such capability, an application server relies on a single reference time source that it cannot validate or check and so trading software timestamps are not verifiable. This is essential even in single site small platforms but becomes even more critical in large scale networks.

Failure Modes of Legacy Technology

This is a summary of problems in the "default" configurations which make use of free software time clients and traditional commercial GPS clocks.

- NIST time and GPS time are not easy to track concurrently and cross-check for sanity.
- GPS/GNSS device failures (lightning on an antenna, cabling, spoofing, GPS/GNSS radio failure or failure to properly interpret current time, networking failures ...). Here are three common types of failure:
    - One widely used GPS network clock device sometimes "forgets" to add in leap second adjustments, causing time to jump backward 35 or so seconds.

    - Many GPS network clocks lack dual power supplies or even proper alarming on hardware failure.
    - Out of date, obsolete, network interfaces can introduce significant “asymmetry” which will defeat the calculations of downstream “client” systems.
- Networking failures, configuration problems, fragile “time enabled” network devices such as switches and routers with built in PTP-1588 support, hidden asymmetries, and even load variation on networks can make time delivery unreliable or even unavailable. These types of errors are common in production systems.
- Widely used free software clients, NTPd and PTPd and variants are unreliable, produce misleading diagnostics, do not have any mechanisms for failover or crosscheck, require development of customized/improvised management software, and have complex, error prone, configuration. In addition, this software has a long history of serious security flaws that can be exploited to defeat time synchronization and to even bring down whole trading platforms.
- The “best master clock” protocol of IEEE-1588 does not address any of the issues of reliability in modern enterprise computing. Best master clock relies on the server computer to inform the client computers which time source has best quality – and the server computer rarely has any information needed to determine time quality at the client. IEEE-1588 Standards committee is, at this time, considering future changes to the standard that might address fault-tolerance, but completing the standard, developing technology around that standard, and validating designs is not likely to be a short term process.

Future issues and cloud

As financial services institutions continue to embrace cloud technology and infrastructures, the ability to achieve millisecond or better levels of precision for time synchronization in the cloud will also become important. Solutions that do not work in cloud environments will need to be supplemented by ones that do.

Conclusion

The proposed standard is a significant step forward but requires proper attention be paid to assurance if it is to be more than symbolic requirement. Meeting the 50 millisecond standard should not be a serious burden for even smaller market participants and it is an essential step towards higher data integrity. Wide use of electronic trading systems and proliferation of trading venues make it impossible to understand market operation or to manage risks without precise and reliable time information. FINRAs proposed rule to implement the Tighten Business Clock Synchronization Requirements is timely and necessary. The technology needed to adhere to the proposed requirement is available to market participants, and the risks are too high to delay implementation of this standard. We commend FINRA for recognizing this important market structure issue.

Thank you for your consideration.

Sincerely,

Victor Yodaiken
CEO
FSMLabs.
www.fsmlabs.com

Appendix: Clock synchronization technology and terminology

NIST: Official time in the USA comes from atomic clocks operated by the National Institute of Standards and Technology (NIST).

GPS: Global Positioning System (GPS) satellites have time from atomic clocks then made available to GPS receivers. In practice, that time is within a few nanoseconds of NIST time. GPS radio signals are weak and can be lost or blocked or spoofed (see below). There are several international GNSS systems that are alternatives to GPS.

Spoofing: Where an attacker attempts to break security by providing false data. GPS spoofing involves overriding GPS signals and sending false timing information that appears to be from the satellites.

Network Clock: A device that receives time from some reference source, generally GPS or NIST servers, and sends that time out over computer network.

NTP: Network Time Protocol is the most widely used network protocol for sending time from a source (such as a device equipped with a GPS receiver) and clients (computers that run application code). NTP is a simple client/server protocol: the client asks the server for the current time, and the server then replies. Part of the complexity is for the client to try to figure out how long the reply has taken to arrive – and to adjust local time with that delay taken into account.

PTP – Precision Time Protocol also known as IEEE-1588 is an alternative to NTP. PTP has multiple modes of operation, many added as it has been modified to work better in enterprise networks. The original design was for very simple networks used in industrial control. PTP has evolved to look a lot more like NTP than it did originally. In PTP terminology, a server is a “grandmaster” and a client is a “slave”.

Hardware Assist: One of the innovations originally associated with PTP was to make network devices assist in computing the delay of time packets. For example, network adapters now sometimes can be configured to tag incoming time packets with the time that the packet arrived at the device. This allows the client software to adjust for time that the packet spent in the operating system network stack. Many network devices now provide this support for both NTP and PTP packets. Additionally, some newer switches and routers can act as in-between servers (boundary clocks) or can add delay information to PTP packets.

Client Synchronization Software: runs on the application server computer and operates an "ideal clock" or "smart clock" that is driven by local timing hardware on the application computer and by information received from reference time sources – such as over the network via NTP or PTP or perhaps a mix of those.