Single source IEEE PTP 1588 cannot meet financial regulatory standards

Victor Yodaiken, 3/30/2016

A 2014 technical paper [IND[i]] written by lead engineers at IMC, NYSE, and Deutsche-Boerse investigates one of the design flaws in IEEE 1588 PTP that makes systems relying on it vulnerable to catastrophic timing errors in ways that would violate financial trading regulatory requirements such as those in MiFID II and CAT[1]. The key point made by the authors is that:

> **the root cause lies in the PTPv2 standard itself: the standard is vulnerable to byzantine failures, so it affects any PTPv2 implementation in which clients trust a single time source**

IND suggests developing solutions that are somewhat similar to the solutions found in TimeKeeper[2], but solutions aside, the paper indicates increasing awareness of "robustness issues" among technically sophisticated financial market firms that had previously relied on the PTP standard.

The problem discussed in the IND paper lies in the "best master clock" algorithm. The PTP standard provides for synchronization of clocks on "slave" devices that receive time over a network from "master clocks". The authors of the standard assumed that master clocks and the time feeds provided by master clocks were only subject to simple failure characterized by the feed turning off. In that case, a second master clock on the same domain could start up and take the place of the first clock. The best master clock algorithm requires slaves to choose the master clock that advertises itself as the most accurate and even requires alternate master clocks to enter a passive mode if they encounter a feed from a master clock advertising higher accuracy. The problem the IND paper identifies is that of time feeds that are compromised but are advertised as accurate. The slave computers which run trading software are required by the PTP standard to unquestioningly accept those feeds even if this causes them to report grossly erroneous times. In fact, as the IND authors point out, these type of clock failures have repeatedly affected financial trading networks in the field. The IND authors write:

> *This paper describes a fundamental single point of failure in the PTPv2 protocol that affects its robustness to failure in specific error scenarios. The architecture design of electing a single unique time source to a PTP domain – the PTP GrandMaster – makes this protocol vulnerable to byzantine failures.*

[1] This is far from the only problem in the standard or in free-software implementations.
[2] The IMC experimental solution outlined in the paper appears to have fundamental robustness and accuracy limitations.

"Byzantine failures", in this case, are failures where the master clock continues to provide a time feed, but provides bad times for one reason or another[3] . The IND paper shows that any firm relying on the ubiquitous free software single-source PTP "clients" (PTPd, PTP4l, SFPTPd, etc) is vulnerable to catastrophic time error resulting from this type of master clock failures. As the authors point out, multi-source operation is essential to coping with malicious attacks as well.

IMC has, at least as a test, added a complex monitoring system outside of the clients in order to mitigate effects of those failures and then developed experimental clients with both PTP and NTP sources that could fall back on an NTP source if a PTP source diverged too much. However, multiple time sources are not sufficient.

> *The test results show that having multiple PTP time sources available is not enough – instead, it is really required that endslaves themselves are able to continuously query and apply offsets from multiple geographically disperse time sources, repeatedly.*

I think that it is absolutely correct that the edge units need multiple sources and intelligent time analysis for fault tolerance. This capability was introduced into TimeKeeper in late 2011 and has been significantly enhanced over 4 years of production use and development. PTP itself is not enough to meet regulatory standards. It must be used in the context of a comprehensive system of quality assurance, fault-tolerance, and monitoring.

> TimeKeeper 5.0 offers the ability to monitor multiple time distribution channels, even those operating on different time distribution standards or of different quality due to distance or network issues. As an example, a TimeKeeper client may monitor two different Precision Time Protocol (PTP) "master clocks" and three different Network Time Protocol (NTP) servers. In addition, if the time quality of TimeKeeper's primary sources becomes questionable, TimeKeeper can now switch from tracking one time source to another, according to a fail-over list provided at configuration time. (FSMLabs Press Release, 2011)

[i] P. V. Estrela, S. Neusüß and W. Owczarek, "Using a multi-source NTP watchdog to increase the robustness of PTPv2 in financial industry networks," *Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS), 2014 IEEE International Symposium on*, Austin, TX, 2014, pp. 87-92.
doi: 10.1109/ISPCS.2014.6948697

[3] I consider "byzantine failure" to be overly cute and imprecise, but it is the standard academic term.