

STP recommendations for the FINRA clock synchronization requirements



This document can be found on the web, www.ibm.com/support/techdocs

Search for document number WP102690 under the category of "White Papers".

Version Date: January 20, 2017

George Kozakos

gkozakos@au1.ibm.com

Contents

Configuring an external time source (ETS) via a network time protocol (NTP) server with pulse per second (PPS)	3
Activating the z/OS time of day (TOD) clock accuracy monitor	4
Maintaining time accuracy across power-on-reset (POR)	5
Specifying leap seconds	6
Appendix	7
References	8
Trademarks	8
Feedback	8
Acknowledgements	8

Overview

New clock synchronization requirements have recently been announced by the Financial Industry Regulatory Authority (FINRA). Beginning February 20, 2017, computer clocks that are used to record events in national market system (NMS) securities and over-the-counter (OTC) equity securities must be synchronized to within a 50 millisecond drift tolerance of the National Institute of Standards and Technology (NIST) atomic clock. This paper documents server time protocol (STP) recommendations to best meet this FINRA clock synchronization requirement.

The following topics will be discussed:

- Configuring an external time source (ETS) via a network time protocol (NTP) server with pulse per second (PPS)
- Activating the z/OS time of day (TOD) clock accuracy monitor
- Maintaining time accuracy across power-on-reset (POR)
- Specifying leap seconds

Configuring an external time source (ETS) via a network time protocol (NTP) server with pulse per second (PPS)

An external time source (ETS) can be used to synchronize the time of the STP-only Coordinated Timing Network (CTN) to Coordinated Universal Time (UTC). An external time source can be configured via an NTP server which can achieve time accuracy within 100 milliseconds. In order to meet the FINRA clock synchronization requirement of 50 milliseconds, the NTP server must have a pulse per second (PPS) output signal that can achieve time accuracy within 10 microseconds.

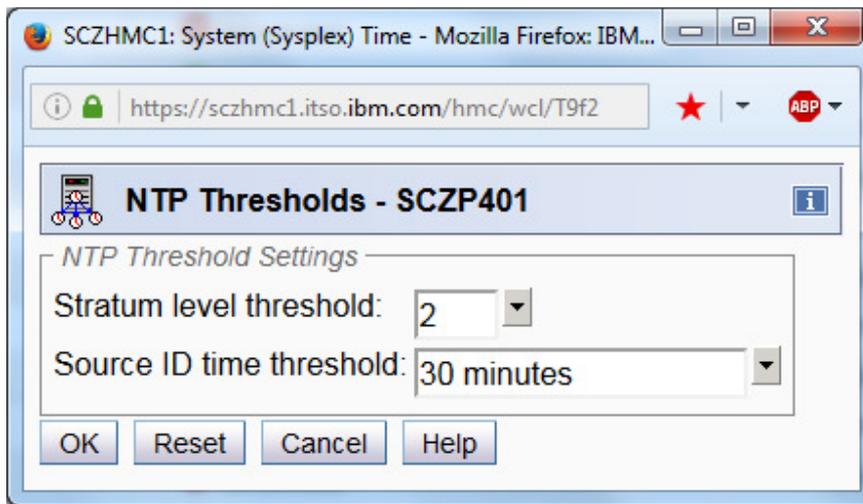
When STP is synchronizing to a recognized stratum 1 time source, a function called fine steering is active. The basis of fine steering is that an oscillator will have a constant error rate. By keeping a history of the adjustments made, fine steering is used to make the expected adjustment if connectivity to the NTP server is lost. The stratum level and source identifier can be seen via the ETS configuration panel on the hardware management console (HMC). IBM recommends that the NTP server is stratum level 1 and the source identifier is one of the known reliable time sources documented in appendix table 1.

The HMC uses NTP version 4 for communication with the selected time source. For information on selecting the NTP time source see:

<http://support.ntp.org/bin/view/Support/SelectingOffsiteNTPServers>

NTP thresholds can be set via the HMC panel shown below to ensure warning messages are issued when the stratum level exceeds stratum 1 or when the source identifier

changes. The recommended settings are a stratum level threshold of 2 and a Source ID time threshold of 30 minutes.



In general, NTP servers have functions to notify the operator if their clock source is lost or compromised. IBM recommends that customers use these functions in addition to other notifications to ensure they remain within FINRA guidelines. The user's manual of the specific NTP server should have details on how to do this.

Activating the z/OS time of day (TOD) clock accuracy monitor

The z/OS TOD clock accuracy monitor was introduced with z/OS V2.1 to monitor the accuracy to which an STP-only CTN is synchronized with an external time source (ETS). The ACCURACY statement in the CLOCKxx PARMLIB member specifies a threshold for the time deviation between the UTC time of the STP-only CTN and an external time source. The accuracy monitor checks the deviation and issues a message if found to exceed the specified ACCURACY threshold.

The z/OS TOD clock accuracy monitor is enabled by specifying a non-zero value for the ACCURACY parameter in the CLOCKxx PARMLIB member. The value must be in the range 0 - 60000. The default value for ACCURACY is zero, meaning that the accuracy monitor is disabled. When CLOCKxx is processed at IPL time, and a non-zero value is specified for the ACCURACY parameter for a system running on an STP network that is synchronizing to an external time source, the ACCURACY function is activated. This results in the following IPL-time message:

```
IEA034I THE TOD CLOCK ACCURACY MONITOR IS ACTIVE
```

The monitor will then check the time deviation every 60 minutes. If the TOD clock exceeds +/- the ACCURACY value, the following message is issued and then reissued every 60 minutes until the condition is corrected:

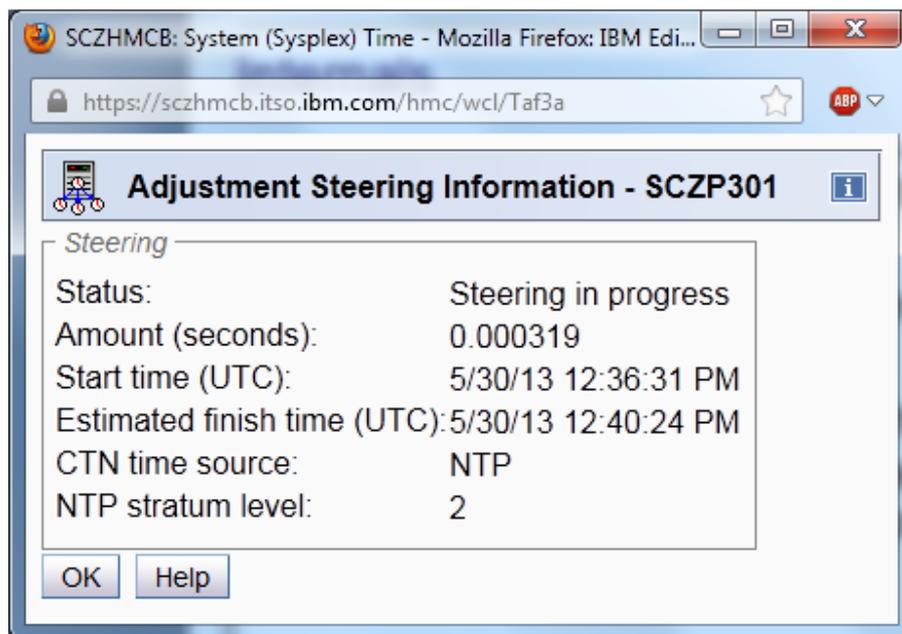
```
IEA032E TOD CLOCK ACCURACY LIMITS MAY HAVE BEEN  
EXCEEDED
```

When the time difference is corrected the following message is issued:

IEA033I THE TOD CLOCK IS NOW WITHIN SPECIFIED ACCURACY
BOUNDS

IBM recommends that ACCURACY be set to 20 for customers that need to meet the FINRA clock synchronization requirement. This will cause IEA032E to be issued when the steering adjustment exceeds 20 milliseconds.

The z/OS TOD clock accuracy monitor uses information from the primary reference time (PRT) correction steering information block. This contains the information seen via the STP Adjustment Steering Information panel shown below.



To help customers meet the FINRA clock synchronization audit requirement, new function will be added to the z/OS accuracy monitor via APAR OA51786. A message will be issued at a user specified interval, with a default of every 60 minutes, giving details of the current status from the PRT correction steering information block. The PTFs for this enhancement will be available by March 31 and a ++APAR by mid February.

Maintaining time accuracy across power-on-reset (POR)

When a server (otherwise referred as a central processor complex (CPC)) in an STP-only CTN is POR'd and then joins the active CTN, the server's TOD clock is jumped to the time of the stratum 1 server. However, when a single server CTN is POR'd or all servers in a multi-server CTN are POR'd, the time for the STP-only CTN is taken from the support element (SE) when the first server is activated. This means the accuracy of the SE's battery operated clock (BOC) is critical when the first server of the STP-only CTN is POR'd.

The accuracy of the SE's battery operated clock (BOC) has been improved significantly in recent years. Enhancements were made with z196 GA2 (driver D93G) to synchronize the SE system clock to the server TOD every hour rather than every 24 hours. In addition the SE system clock is now steered rather than jumped to ensure there are no duplicate timestamps, and the SE BOC is updated from the SE system time. This means that while the STP-only CTN is synchronized to an external time source, the SE system time will be accurate, but the SE system time will drift when the CTN is down.

There were further enhancements with zEC12 GA2 (driver D15F), which was released in September 2013:

- For STP customers, when the server is down, the SE time will be updated every hour via NTP server, but only when an NTP server is configured via the STP ETS configuration panel.
- When the SE is IMLed, if the server is down, the SE will access the ETS to get the time. It will continue to get the time from the ETS source every hour while the server is down. If the server is up, the SE will get the time from the server.
- When the server is POR'd, the SE will get the time from the ETS before the server's time is set.

Specifying leap seconds

When a leap second is inserted into Co-ordinated Universal Time (UTC), UTC effectively jumps backward by one second. If leap seconds are not specified for the STP-only CTN, UTC will be one second ahead of the NTP server that STP is synchronizing to when the leap second event occurs. STP will automatically steer to the correct time but this steering adjustment is made at the rate of seven hours per second to ensure no duplicate timestamps are seen.

In order to maintain time accuracy when a leap second event occurs, leap seconds must be specified and any leap second adjustment must be scheduled for the appropriate date and time via the STP Adjust Leap Seconds panel.

Co-ordinated Universal Time (UTC) was introduced in 1972 and was 10 seconds behind international atomic time (TAI). Leap seconds are used to synchronize UTC with TAI and there have been 27 leap seconds since 1972. The system z clock is architected to be 10 seconds behind TAI and so the current leap second value is +27 seconds, with the most recent leap second having been added at midnight UTC on December 31, 2016. When a leap second is needed, it will generally be added at midnight UTC on either June 30 or December 31.

Customers who do not specify leap seconds will not meet the FINRA clock synchronization requirement during the seven-hour window required for steering to complete. This will be highlighted by message IEA032E issued by the z/OS accuracy monitor at every 60 minute interval until steering has corrected the time to within the ACCURACY threshold.

Customers that plan to start specifying leap seconds to ensure time accuracy when leap seconds are scheduled have two choices:

- The leap second value can be updated from 0 to +27 and then kept up to date when new leap seconds are scheduled. As a positive leap second change is a negative UTC change, z/OS spins for the amount of the positive leap second change to avoid duplicate time stamps. A spin of 27 seconds will very likely cause a system outage and so leap seconds need to be added in small increments or while the z/OS systems are down.
- Alternatively, leap seconds can be added starting from when the next leap second is scheduled. When the next leap second is added, the STP leap second value can be updated from 0 to 1 and then incremented as each subsequent leap second is added. This will ensure time accuracy across future leap seconds without the problem of getting from a current value of 0 to 27.

Appendix

Source Identifier	External Reference Source	Dispersion
LOCL	Uncalibrated local clock used as a primary reference for a subnet without external means of synchronization	not applicable
PPS	Atomic clock or other pulse-per-second source individually calibrated to unknown national standards	
ACTS	NIST dialup modem service	15 milliseconds
USNO	U.S. Naval Observatory modem service	15 milliseconds
PTB	PTB (Germany) modem service	15 milliseconds
TDF	Allouis (France) Radio 164 kHz (LF)	15 milliseconds
DCF	Mainflingen (Germany) Radio 77.5 kHz (LF)	15 milliseconds
MSF	Rugby (UK) Radio 60 kHz (LF)	15 milliseconds

WWV	Ft. Collins (US) Radio 2.5, 5, 10, 15, 20 MHz (HF)	20 milliseconds
WWVB	Boulder (US) Radio 60 kHz (LF)	15 milliseconds
WWVH	Kauai Hawaii (US) Radio 2.5, 5, 10, 15 MHz (HF)	20 milliseconds
CHU	Ottawa (Canada) Radio 3330, 7335, 14670 kHz (HF)	20 milliseconds
LORC	LORAN-C radio navigation system (LF)	15 milliseconds
GPS	Global Positioning Service	150 nanoseconds
CESM	Calibrated Cesium clock	unknown
RBDM	Calibrated Rubidium clock	unknown
FLY	Source is flywheeling on its local oscillator after having achieved synchronization	unknown

Table 1: NTP Stratum 1 clock source identifiers and dispersions

References

FINRA Regulatory Notice 16-23:

http://www.finra.org/sites/default/files/notice_doc_file_ref/Regulatory-Notice-16-23.pdf

Trademarks

You can find a full list of U.S. trademarks owned by IBM at:

<http://www.ibm.com/legal/us/en/copytrade.shtml>

Feedback

Please send comments or suggestions for changes to gkozakos@au1.ibm.com.

Acknowledgements

The author would like to thank the STP design team with special thanks to Donald Crabtree, Steve Cherniak and Peter Relson.