# MiFID II TIME SYNCHRONISATION COMPLIANCE POSITION PAPER

16 MAY 2016

## Overview and purpose

This position paper describes the methodology that Bats Europe is testing and plans to use to demonstrate time compliance to MiFID II as per RTS 25. We believe the configuration described is sufficient to demonstrate traceability to UTC within the 100us requirement, and is robust enough to provide a reasonable level of protection against most GPS anomalies and issues.

We want to emphasise that even though this solution could work for Bats Europe, it might not be right for everyone. As a trading venue we benefit from having all our production infrastructure in the same location and so do not have to contend with the added complexity of enforcing time synchronisation between multiple, geographically-diverse locations.

Our main goal is to implement a simple yet robust solution that can reliably distribute accurate time to our trading environments. Multiple technologies described here are in testing or will soon be tested by Bats Europe to ensure they meet our needs. As such, it's possible that the design will need to be updated depending the outcomes.

## Resilient time sources and time comparison

Under this proposal, each of our trading environments will be configured as described below. Each host subscribes to an active PTP feed from each local PTP Grandmaster (GM) and each GM operates on a different PTP domain. All GM PTP feeds are active at all times. One is the primary, and the second provides resiliency and an additional source against which the primary can be compared to detect time drift.

Time is distributed to the hosts reliably over a simple one-tier 10G network which helps to ensure minimal jitter for delivery of PTP messages. This is the same network the hosts use for their backend communications and is not dedicated to PTP messages. Each host consuming the PTP feed also uses a local hardware oscillator to provide limited holdover capabilities in a worst-case scenario.

An option is to use three local PTP GMs. With this option, each generic host can compare three PTP times to each other and continually select the best. When two PTP GMs are used, this selection cannot happen automatically. However, the Time Comparison Host will alert on any drift in PTP GM1 or GM2. Following that, manual intervention can be taken to run the environment's hosts on PTP GM2 if GM1 is out of sync. We are investigating whether this option is worth implementing.

There is a level of monitoring on each host to ensure the hosts don't drift out of sync or take in erroneous time data via PTP. To achieve this, the deltas between the feeds and the host's clocks are continually logged and alerts generated if they drift out of sync. In a worst-case scenario the host clock can be manually reconfigured to run from its local oscillator until the error is corrected.

**Diagram 1**

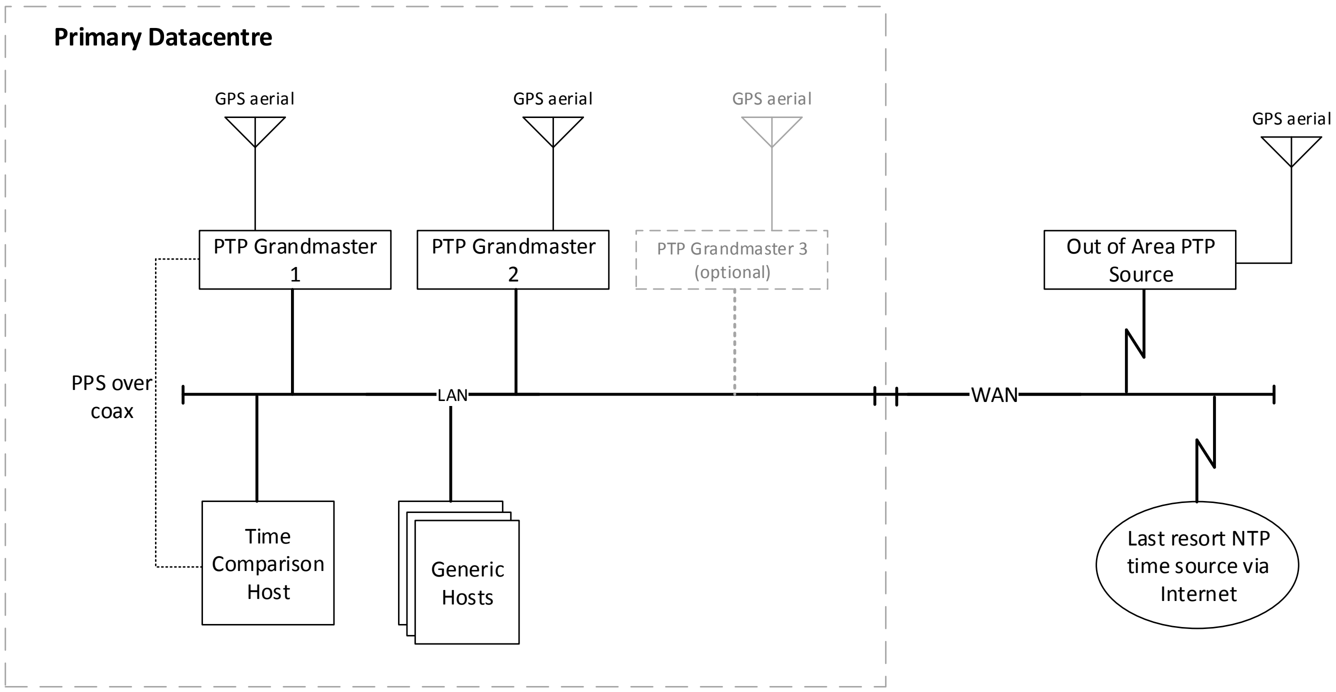


In addition to the design shown in diagram 1, each trading environment has a "Time Comparison Host". This host continually pulls in time information from multiple sources, compares them to each other and records all deltas for alerting and graphing.

The sources proposed are as follows;

1. Local PTP GM 1 - Primary Datacentre, Multicast PTP
2. Local PTP GM 1 - Primary Datacentre, PPS over coax
3. Local PTP GM 2 - Primary Datacentre, Multicast PTP
4. Local PTP GM 3 - Optional (automatic selection of best out of three)
5. Remote PTP GM - Out of Area Datacentre, Unicast PTP
6. Remote NTP - Public NTP Service, Internet

A further description of each source and its purpose is as follows:

1. The local site primary PTP source. The hosts run with this as the active time source under normal operating conditions.

2. A PPS feed from the primary PTP source to the time comparison host serves to establish the delay inherent in the local delivery network. This will also enable us to react to fluctuations in delivery of PTP over the network, and provide appropriate alerts.

3. The second local PTP GM serves as a redundant time source. This is completely separate hardware from the primary, and as such another source of time against which to compare.

4. An optional third local PTP source which will enable all PTP-enabled hosts to automatically pick the best of three.

5. An out-of-area source providing PTP. This would help show GPS-related problems that affect our primary site, though not a larger area.

6. As a last measure we would take a reliable public NTP feed. As NTP is provided via the Internet it is not up to microsecond accuracy. This feed can potentially highlight major issues with the global GPS system.

## Demonstrating and documenting compliance

As a solution based on solid vendor specifications and with a high level of monitoring via the time comparison hosts, we believe Bats will be able to demonstrate compliance with data showing deltas between each of the time feeds seen by the time comparison host

Bats also supports the idea of a standard form describing each component of the time distribution chain, including offset, uncertainty and monitoring for each component. For some components, that information will come from the manufacturer's specifications and from practical testing of the implemented solution.

# PTP Distribution over a local area network

The time comparison functionality we have described can be used mostly irrespective of how PTP time is actually distributed to the end hosts. Bats cannot discuss its network layout in any detail. We can say that the basic principle of PTP time distribution has to be to keep the network layout as simple as possible. The PTP GM should ultimately be sending PTP packets to the hosts over the shortest path possible. This will reduce any jitter introduced by more complex network layouts.

Diagrams 2 and 3 show two examples of how PTP time can be distributed in a simple way to multiple hosts and multiple subnets within a datacentre. This is by no means a complete list of possible architectures but rather just two basic examples on which to build. We also don't show any resiliency options but that could be addressed in a number of ways, such as duplication of the topology for a completely redundant path to hosts with multiple interfaces.

Diagram 2 shows a PTP GM with multiple Ethernet interfaces. In this case each host subnet connects to a dedicated physical Ethernet interface on the GM, keeping the path between the GM and hosts as simple as possible. The number of subnets that can be supported is limited by the number of physical Ethernet interfaces on the GM.

**Diagram 2**

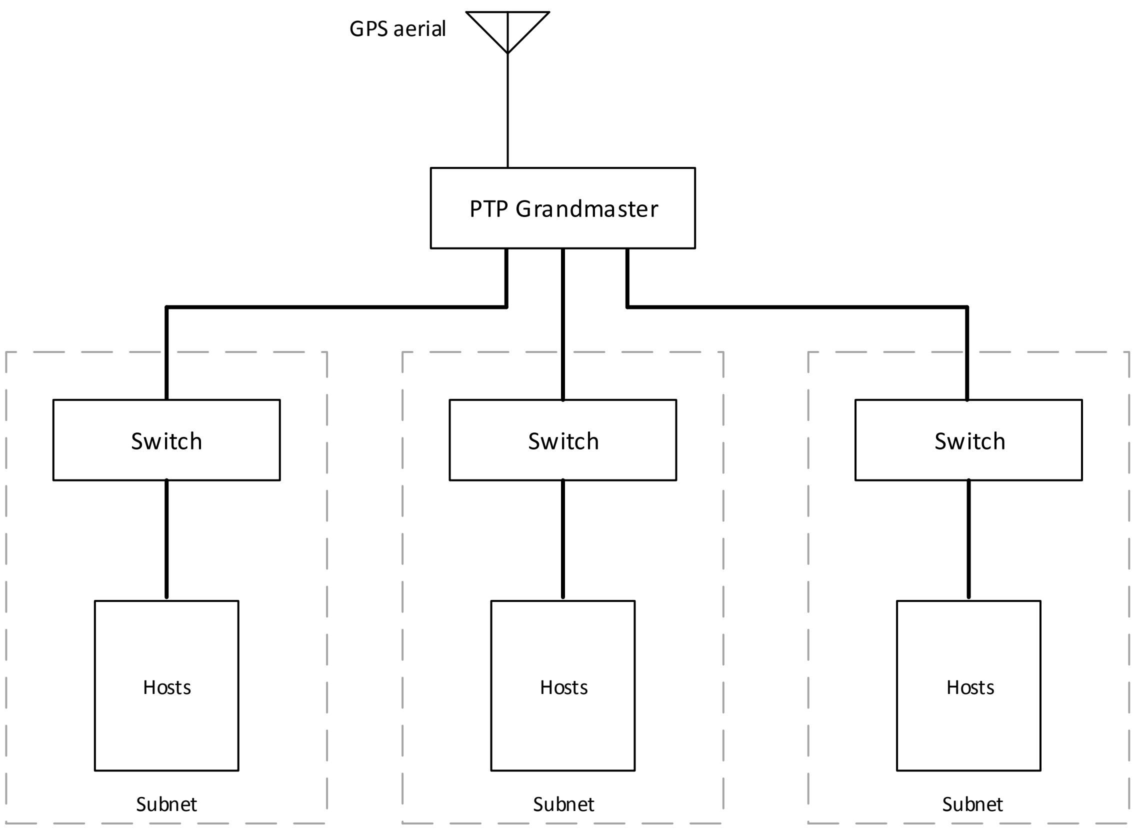

Diagram 3 shows a PTP GM that distributes PTP to multiple boundary clock PTP appliances, each of those is servicing hosts in one subnet.

This method is more complicated as it requires multiple PTP devices but can scale much bigger. We try to keep the network as simple as possible with only one extra switch hop over diagram 2.

**Diagram 3**

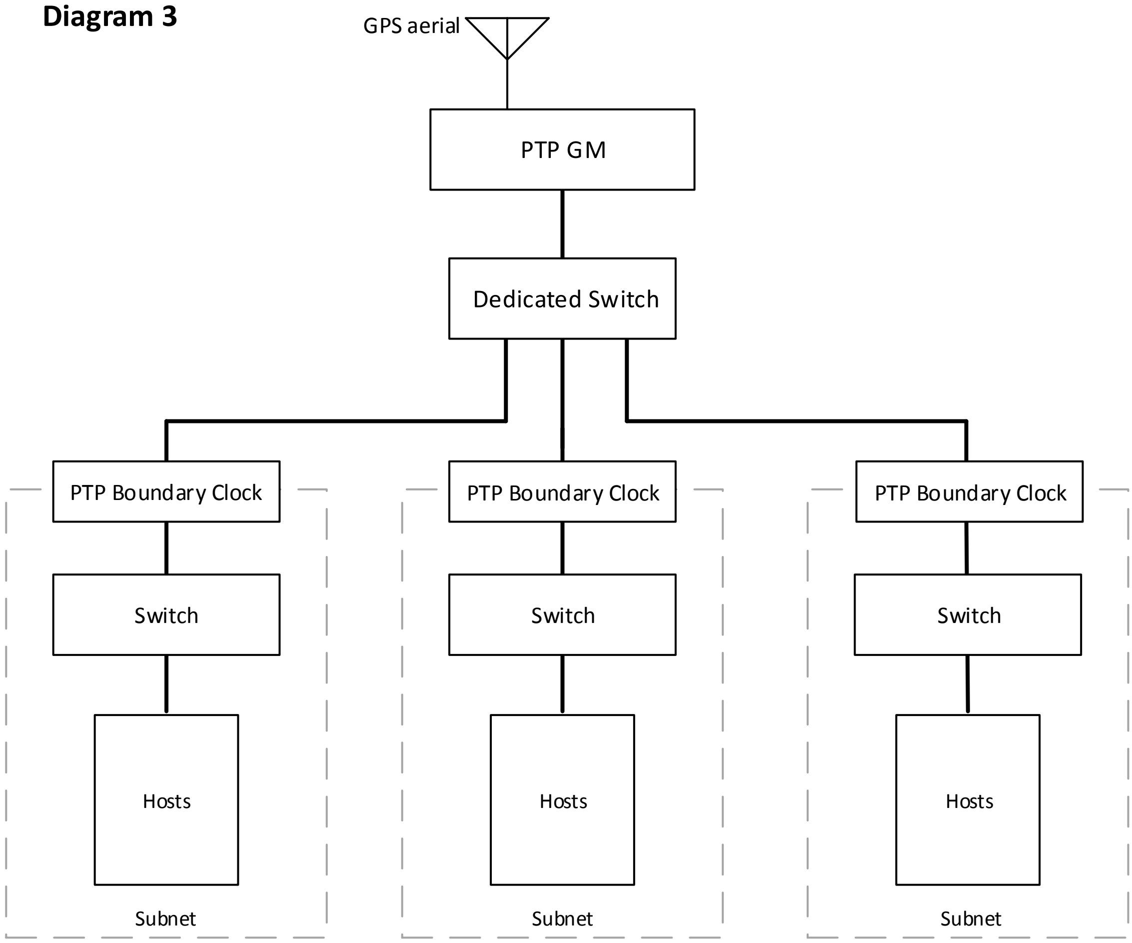


The above designs show two possible ways to address the challenges inherent in distributing time to a large number of hosts on multiple networks with the least amount of jitter possible.

Using a single interface on a PTP GM connected to a multi-tier network can introduce more jitter and more points of potential issues, for this reason we want to keep the PTP boundary clock as close to the host as possible. Where possible jitter should also be continually monitored by each time comparison host by comparing PTP delivery over the network to its PPS feed from the same PTP clock, as shown in diagram 1.

WAN networks will present different challenges and other solutions may need to be considered.

## Summary

Some technologies described in this document are outside the current IEEE 1588 PTP standard. PTP clients actively listening to multiple PTP domains and selecting the best is an example of such a feature, it's not in the standard but is a vendor addition. Bats will evaluate and make use of such features if they provide useful enhancements with no detriment to accuracy. All such features described in this paper are subject to further verification and testing by Bats.

By adopting a simple network layout that works for our situation we can keep complexity in delivery of PTP messages to a minimum and focus our efforts on robust delivery and comparisons, logging and alerting on any time drift detected by our comparison hosts.

In the proposed solution the installation and configuration of the GPS aerials, cables and PTP GM's will follow established industry best practices and manufacture guidelines. Bats Europe works closely with the manufacturers of each component of the solution to ensure we fully meet their installation specifications.

Bats will update this document as we progress further with our testing of the proposed solution and validation of the functionality of all its components.

This document is available on the Bats Europe MiFID II microsite at
http://www.bats.com/europe/equities/regulation/mifid/