TECHNICAL BRIEF

One of the most prevalent protocols for industrial time transfer is Precision Time Protocol (PTP). PTP offers dependable, high-accuracy time synchronization of multiple clocks over Ethernet networks enabling real-time systems and applications to operate correctly. This document provides an overview of the benefits and challenges when choosing a PTP Profile for a particular application, discusses appropriate hardware and software capabilities for network equipment, and examines the impact of different network environments on time accuracy.

Calnex
Insight and Innovation

Time Transfer in Industrial Ethernet Applications

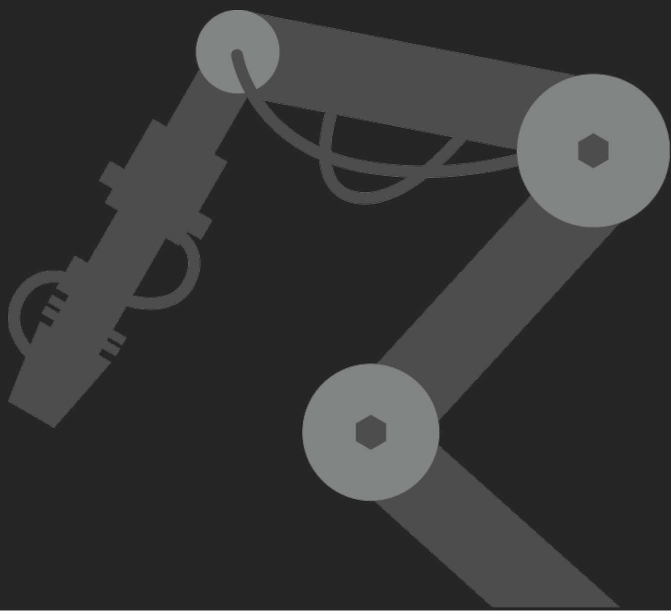

Time Transfer in Industrial Ethernet Applications

Within several industries, Precision Time Protocol (PTP, defined in IEEE1588-2008) is the preferred method of transferring time for applications with varying levels of accuracy requirements. PTP is a mature solution with industry-specific profiles, special hardware, and rules for network topologies – all of which are potentially required to ensure time-sensitive applications meet strict performance requirements.

Today's Telecoms Industry has deployed networks with the objective of having an end-to-end time accuracy of 1µsec, utilizing both hardware timestamping and timing support in some or all of the devices used in the network. Industrial Ethernet applications may, on the surface, have a less stringent accuracy requirement compared to telecoms, but that does not mean the challenge of accurately transferring time is by any means less. While Telecom's applications require Time to be accurate to a few microseconds, the environments and system requirements in Industrial Ethernet applications are very different such that delivering time accurate to just hundreds of milliseconds may prove equally challenging.

Selecting a Profile

IEEE1588-2008 is an umbrella document defining all the parameters and constructs that may be required when transferring time using PTP. The expectation is that each industry produces a Profile that defines the sub-set of features and parameter values that are appropriate to that specific industry application.

For example, the Power Industry has defined a PTP Profile listing the specific sub-set of features and appropriate parameter values for use in power station networks. This is documented in IEEE C37.238. Similarly, the audio/video industry has defined a PTP Profile originally aimed at consumer audio/video networks in IEEE802.1AS. This is now being extended and adopted as part of IEEE's "Time Sensitive Networking" program, and is targeted at applications such as in-car networks, factory and industrial networks, and the "Internet of Things".

One of the benefits of having a defined Profile is that interoperability between networking devices conforming to that Profile can be assured plus, in most cases, performance specifications will already be specified. The latter point is important as PTP does not define or set any performance requirements, rather it is up to the Profile to set performance requirements appropriate for the target application. Likewise, network equipment conforming to the IEEE1588-2008 standard does not guarantee interoperability either as each vendor my select a different sub-set of the requirements to align with.

The ITU-T has also defined Profiles for the transfer of Time, namely G.8275.1 and G.8275.2. While these are defined for the Telecoms Industry, there is a wide range of equipment available that conforms to these standards. As stated above, it is better to have all equipment conforming to the same Profile than none at all.

Sources of Inaccuracy

In the generic network shown below, there are three sources of inaccuracy, namely the Grand Master clock reference, the Ethernet switches, and the PTP slave/end device inside the controller within the end-equipment rack.

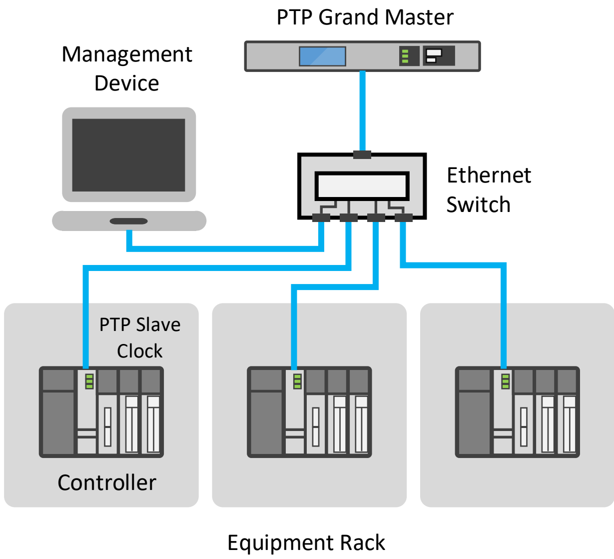

Grand Master

In large production facilities there may be many tens, if not hundreds, of controllers that all need to have accurate time delivered to them. This is a high demand and one that is best handled by using a dedicated PTP Grand Master (GM). An off-the-shelf GM is likely to deliver Time Error performance in the region of 100ns which, for most applications, will be well within the overall accuracy requirements.

However, if the accuracy delivered by the GM is far more than required by the application, this may beg the question 'why not make one of the controllers the GM?' (that is, implementing the GM function in software). The problem here is that without hardware support it is hard to envision any software system could guarantee its performance to all controllers given there is no synchronization between when they generate messages towards the GM. Messages from multiple end points arriving concurrently could induce significant delays in the processing time to action each one. And many messages arriving simultaneously could lead to severe congestion in the GM function resulting in high inaccuracy unless the GM is specifically designed to handle this situation with defined accuracy.

With an architecture that uses a single (or small number of) GM to provide time for the whole production line, the cost of buying a single GM that will guarantee its time accuracy is highly likely to outweigh the cost of developing and deploying a custom built GM which will have all the challenges in controlling the inaccuracy under all scenarios.

Interconnection Network

Unless a separate network is deployed purely for the PTP flows, which is an expensive option, the PTP flows will have to travel through the same network as all the data traffic. Each flow might only pass through a small number of switches, however, as it will be carried through switches and links with other traffic, there is the possibility that the timing packets may experience a wide range of transit delays.

It should be noted that the absolute delay through the network is not the challenge with PTP as this can be measured and hence compensated for. The challenge is packet-to-packet variation which impacts the accuracy of the time required in the end device. In addition, asymmetry through switches (the up-stream transit delay being different to the down-stream transit delay) also produces inaccuracies in the recovered time because the PTP algorithm is blind to asymmetric delay. PTP measures the round-trip delay and divides by two to estimate the one-way delay, in other words, it assumes symmetry.

So what actions can be taken reduce the delay variation?

First, configuring the switches such that PTP flows are assigned the highest priority. This should help to reduce the delay variation experienced by the PTP flows, but it will not remove the variation as this will depend on whether any other traffic flows also need to be given the highest priority.

Second, by avoiding the use of jumbo or long packets in the network. Even when a PTP packet is assigned the highest priority, if it arrives at the queuing buffer that has just started the transmission of a jumbo packet, the PTP packet has to wait until the whole of the jumbo packet has been transmitted before its transfer can start. If it is possible to manage the sizes of the packets in the network, this can help to reduce the delay variation.

In addition, depending on the topology of the network, especially if there are a small number of switches, it is worth considering purchasing switches that are PTP aware – switches that identify PTP packets and manage their transfer through the switch using the Boundary Clock or Transparent Clock function (refer to 'Time and Time Error' on the Calnex website for descriptions of these devices). Telecom devices will deliver time accuracy in the sub 100ns range which is likely to be more than sufficient for this application.

If PTP-aware switches are not used, it is important to develop appropriate test scenarios to ensure the delay variation produced by the variation in network traffic does not impact the ability to recover time in the end devices. Section Proving Performance below discusses how to develop these testing scenarios.

End Device, PTP Slave

Production lines typically have many end points hence many instances of PTP Slaves. Buying a separate termination device is a straightforward solution, however, this will increase the equipment spend and may lead to problems locating the device depending on the physical structure of the end point rack.

Another solution to consider would be to implement a PTP Slave within each rack controller. Developing it purely in software offers an 'easy' deployment route as it does not require any physical change to the controller or rack, only a software update. However, irrespective of the accuracy required, it should not be assumed the target accuracy can be achieved. Complex software systems are notorious for timing inaccuracies in PTP Slave implementations because of the variation in time from when a PTP packet arrives until the point when it is actually processed. Many factors, such as parsing the PTP packet, the time taken to schedule the PTP packet handler, etc., affect the time this will take (remember, it is the variation that causes inaccuracy in the recovered time) therefore careful design analysis coupled with effective test scenarios are essential to ensure the required accuracy is achieved under all operating conditions.

It should be noted that simple averaging of the measured round trip delays is unlikely to deliver the accuracy required, especially if significant delay variation is expected. It is typical that packet selection techniques look for packets near the floor delay (packets that have not been delayed in any queuing buffer) and discard the other packets. (Refer to 'Time and Time error' on the Calnex website for a discussion on causes and sources Time Error.) ITU-T Standard G.8260 Appendix I discusses packet selection for metrics which provides insight to how the Telecommunication Industry has approached this problem.

In multi-process systems, the challenge is to determine what the maximum delay the PTP process will experience from when it requests processing time until it is granted access to the processing engine. In addition, when the process is granted access, will the process be allowed to run to completion or can it be interrupted by another process before it completes its processing of the received packet? Defining these parameters is critical to determining whether a software only implementation is a viable option to meet system requirements.

Proving Performance

It is vital that performance tests are executed on the system to ensure the timing accuracy requirements are achieved under all conditions.

Note that prior to running any performance tests, a method of monitoring Time in the end device must be established, otherwise it may be impossible to know what accuracy is being achieved across the system. In Telecoms, a 1PPS output is usually provided from the end device's internal clock (interface defined in ITU-T G.8703-2016). This output produces a pulse once per second, with the leading edge of the pulse aligned to the 'Top of Second' in the device (refer to 'Time Error Results' (Doc. CX5010) on the Calnex website for further information). This, therefore, can be used for external monitoring of timing accuracy.

There are two system dynamics that need to be stressed to prove timing performance;

1. **Congestion in the packet network:** Careful thought needs to be given to traffic congestion that may exist in the packet network and hence impede the transfer of timing packets from the GM to the end devices. Each time a network scenario is run, the actual delay variation experienced by the PTP packets will vary as it is dependent on the precise packet-to-packet relationships. It is therefore recommended that the delay variation is recorded and replayed to stimulate the end device to determine its robustness to network effects.

   This can be achieved by using a device with an embedded PTP Slave running a PTP active session to the system's GM, or by using a device that can monitor PTP packets passing between the GM and the end device. Then, by stimulating the end device with a PTP flow impaired by this recorded delay pattern, it is possible to assess, in a controlled environment, the robustness and hence the time accuracy possible in the presence of network congestion.

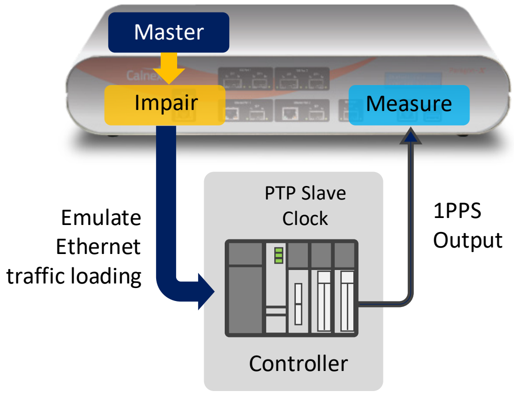

2. **Impact of the multi-process software environment**: If the PTP algorithm is implemented in software then it is vital the impact of the multi-process software environment is verified. The analysis of the software environment (outlined above) during the design phase can be used to specify a number of scenarios which create varying levels of stress on the software system to determine its impact on time accuracy.

   Starting with a scenario where as many of the other processes as possible are disabled, idle or near-idle state is recommended as a baseline. Then, a number of scenarios should be run (perhaps four or five but it depends on the number and variation of loading that may exist in the software environment) where the software system is heavily loaded and the PTP process execution time is impacted. This will determine the worse-case Time Error produced by the software environment.

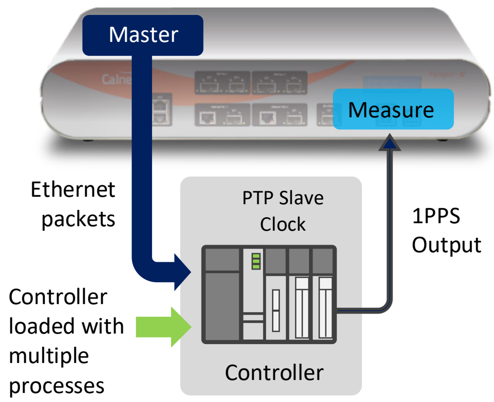

The maximum Time Error for the system is the combination of the worse-case error of the two scenarios.

Summary

PTP is a very versatile and effective protocol for transferring time across packet-based devices and networks. There are already many examples across multiple industries that have used and gained experience on how to use this protocol to achieve the time accuracy required for their applications. Each application will bring different challenges and compromises in order to meet the specific deployment and accuracy objectives.

What is clear from all applications is that performance always requires careful thought and consideration, irrespective of the absolute accuracy required. It has already been observed that very significant inaccuracies can be produced where no consideration of the dynamics of the system are taken into account. There are no absolute must do’s and don’ts, rather it all depends on the specific requirements and objectives of the specific application. However, by utilising the knowledge that already exists, consideration during the design phase of the system dynamics that can affect performance, and development of an effective verification plan, it should be possible to use PTP to achieve the performance goals your system and applications need.

Calnex Solutions Ltd
Oracle Campus
Linlithgow
West Lothian EH49 7LR
United Kingdom

tel: +44 (0) 1506 671 416
email: info@calnexsol.com

calnexsol.com