# IEEE 1588 Precise Time Protocol:
The new standard in time synchronization
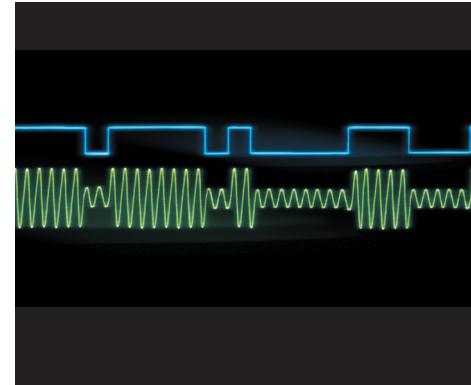
Overview – Precise Time and the Trouble with Ethernet

As network computing becomes more complex and the world more interconnected, the need for more precise time synchronization has greatly increased. Network time protocol (NTP) and IRIG time code have been the most common protocols governing time transfer, but the recently developed IEEE 1588 Precise Time Protocol (PTP) now promises to revolutionize time synchronization by improving accuracy and reducing cost. While certain other precise sync protocols require significant investment in hardware and cabling, PTP makes highly precise timekeeping possible using the most widely deployed medium for network connectivity – Ethernet.

While Ethernet has proven to be a ubiquitous and inexpensive medium for connectivity, it has not been well-suited for applications requiring precise synchronization. By nature it is nondeterministic, which creates difficulty for real-time or time sensitive applications that require synchronization. On an NTP-based LAN, network devices and computer operating systems add latency and jitter that reduces synchronization accuracy to 1 to 2 milliseconds. Applications requiring greater accuracy have often had to deploy separate cabling systems and dedicated clocks – the IRIG B protocol, for instance, requires a dedicated system of coaxial cables to carry timing signals directly between IRIG B clocks separate from any data network.

PTP overcomes the Ethernet latency and jitter issues through hardware time stamping at the physical layer of the network. The result can be an unprecedented accuracy in the 10 to 100 nanosecond range that is achieved using an Ethernet network to carry the timing packets, allowing for remarkable cost savings. This paper examines how PTP works, compares its effectiveness to other protocols, and explains why it is so cost effective.

Choosing a Standard: NTP vs. IRIG vs. PTP

NTP has been the most common and arguably the most popular synchronization solution because it performs well over LANs and WANs and is relatively inexpensive to implement, requiring little in the way of hardware. And while NTP should be able to deliver accuracy of 1 –

Protocol	Sync Accuracy	Interconnect	Required Clock H/W & S/W
NTP	1-10 milliseconds	Ethernet LAN or WAN	H/W or S/W server, S/W clients
IRIG	1-10 microseconds	Coaxial Cable	H/W master and slaves
IEEE 1588	20-100 nanoseconds	Ethernet LAN	H/W master and slaves

Figure 1 Comparison of the synchronization requirements of NTP, IRIG time code and IEEE 1588

2 milliseconds on a LAN or 1 – 20 milliseconds on a WAN, this is far from guaranteed network-wide. This is largely due to the use of switches and routers on LANs and WANs and the fact that many NTP clients run on non-real-time operating systems such as Windows or Linux, that were not designed for time keeping accuracy and therefore rarely deliver it. On Windows, for instance, one can often witness clock corrections of 10 – 50 milliseconds because the system was busy performing tasks it deemed more important than timekeeping.

IRIG time code provides increased accuracy – up to 1–10 microseconds – and is often used where precision timing is mission critical: in military, aerospace and power utility instrumentation. But improved accuracy comes at a cost. IRIG systems eschew Ethernet in favor of dedicated coaxial timing cabling between dedicated hardware clocks, which are not only an added expense but which create an extra burden on the physical infrastructure of the facility.

PTP, on the other hand, offers the cost-effectiveness of NTP by using existing Ethernet LANs, and it exceeds the accuracy of the IRIG clocks. PTP can coexist with normal network traffic on a standard Ethernet LAN using regular hubs

and switches, and yet provides synchronization accuracy to the sub microsecond level. With the addition of IEEE 1588 boundary clocks or transparent switches, 20-100 nanosecond synchronization accuracy is achievable. The key to this caliber of performance is hardware-assisted time stamping.

How PTP Works: Hardware-Assisted Time Stamping

The two primary problems that must be overcome in network timekeeping are oscillator drift and time transfer latency. Regardless of the protocol used, oscillator drift can be mitigated by using higher quality oscillators and by deriving time from a more accurate source, such as GPS. The time transfer latency problem is more difficult and is two-fold in nature: there is latency associated with processing of time packets by the operating system, and network latency created by the hubs, switches, cables and other hardware that exist between clocks. It is in the area of reducing operating system latency that PTP is most successful.

PTP combines time stamping units (TSU) with an innovative method for exchanging time stamp detail between master and slave clocks. A TSU placed between the Ethernet Media Access Control (MAC) and the Ethernet PHY transceiver sniffs both inbound and outbound traffic and issues a time stamp when the leading bits of a 1588 packet are identified, precisely marking the arrival or departure of 1588 time packets.

In order to estimate and mitigate operating system latency, the master clock periodically sends a *Sync* message based on its local clock to a slave clock on the network. The TSU marks the exact time the *Sync* message is sent, and a *Follow_Up* message with the exact time information is immediately sent to the slave clock. The slave clock time stamps the arrival of the *Sync* message, compares the arrival time to the departure time provided in the *Follow_Up* and is then able to identify the amount of latency in the operating system and adjust its clock accordingly.
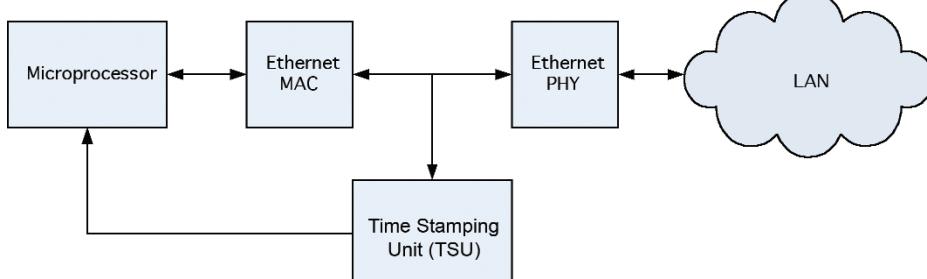
Network related latency is reduced by measuring the roundtrip delay between master and slave clock. The slave periodically sends a delay request message (*Delay_Req*) to the master clock which issues a delay response message (*Delay_Resp*). Since both messages are precisely time-stamped, the slave clock can combine this information with the detail from the *Sync* and *Follow_Up* messages to gauge and adjust for network-induced latency. The protocol for exchanging precise time stamps is detailed in the graphic below.
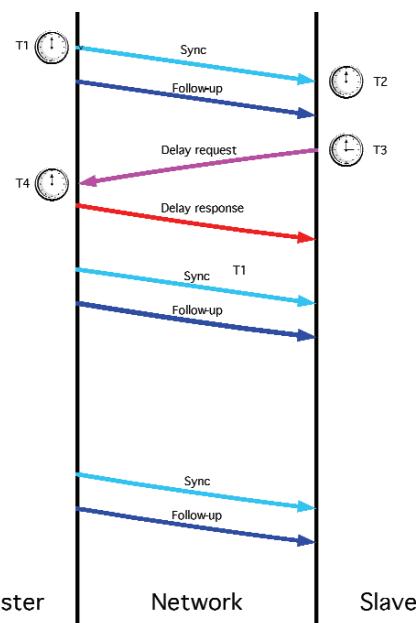
Figure 3 The sequence of packets used to transfer time from the 1588 master to a 1588 slave. Sync packets are stamped as they leave the master (T1) and arrive at the slave (T2). Follow-up packets communicate the sync packet departure time to the slave. Delay response packets are stamped as they leave the slave (T3) and when they arrive at the master (T4). Sync and follow-up packet pairs are sent by the master on a periodic basis, as are the delay request and delay response packets. The formula for slave clock correction= 0.5 (T1-T2-T3+T4).

Determining Target Accuracy

Choosing broadcast intervals and oscillator types

In PTP, the desired accuracy of timing determines how often sync messages are broadcast and what kind of oscillator is used. More frequent broadcasts result in more accurate sync, but also in more network traffic, although the bandwidth used is extremely small. Higher quality oscillators also result in more accurate sync. It may be tempting to try to achieve target accuracy more economically by using a lower quality oscillator while increasing broadcast frequency, but this is unadvisable. Low quality oscillators lack the stability needed to achieve high precision with PTP, so shortening the broadcast interval offers diminishing returns.

Accuracy is also a function of the IEEE 1588 master clock, called the grandmaster, that is the ultimate source of time on the network. Grandmasters are typically referenced to GPS so that they are both very stable and very accurate. Accuracy to UTC, TAI or GPS is typically 30 nanoseconds RMS or better. By starting with such an accurate clock with an absolute time reference, time on a PTP-enabled network can be very well synchronized. A quality grandmaster also provides other measurement features to characterize the latency and jitter characteristics of network elements and to measure slave accuracy relative to the grandmaster. Ironically, a quality grandmaster clock also supports NTP and IRIG timing requirements. This makes sense as there are legacy systems in place that will need synchronization as well as new PTP deployments.

Selecting other hardware

PTP has been readily implemented to work on Ethernet LANs where router buffer delay and switch latencies undermine the accuracy of time transfer. Consequently, the other variables in achieving target time transfer accuracy are the use of hubs versus switches, boundary clocks or transparent clocks.

When evaluating cost versus packet delay, hubs are generally considered the best value. Switches may move packets more efficiently, but reduce accuracy because of packet jitter.

Figure 2 Hardware set up for a system deploying IEEE 1588. The time stamp unit monitors packets at the MII layer, between the Ethernet MAC and PHY, and issues a precise time stamp when a 1588 timing packet passes through it.

The boundary clock is a multi-port switch containing one port that is a PTP slave to a master clock, while the other ports are masters to downstream slave clocks. Boundary clocks provide a decent method for regulating time to a number of subnets. But using cascading boundary clocks accumulates non-linear time offsets in their servo loops, resulting in unacceptable degradation of accuracy.

The transparent clock is another potential hardware option for the PTP-based network. This is a PTP enhanced switch which modifies the precise time stamps in the Delay_Resp and Follow-Up messages to account for receive and transmit delays within the switch itself. The result is improved sync between slave and master clocks. But the transparent clock can also create security issues when the original packet crypto checksum doesn't match the final packet arriving at the slave.

The graphic below shows overlaid histograms of the clock-offset errors relative to each device type. Inconsistent device latency and jitter results in increased clock offset errors represented by the increased dispersion in the data.

Enabling Cost-Effective Real-time Applications

In addition to extraordinary accuracy, PTP also facilitates the operation of real-time or near real-time data acquisition systems over Ethernet. Traditionally, these systems have depended upon triggering techniques and delay compensation independent of the data network in order to coordinate synchronized data collection. But PTP eliminates the necessity – and the cost – of such measures.

Consider the example of a real-time control system that is set up over Ethernet with a variety of independent sensors. Each sensor uses an onboard 1588 clock slaved to a grandmaster clock and is synchronized with the other sensors on the network to an accuracy of 1 microsecond. Every 5 milliseconds, each sensor takes a measurement and sends the information back to the controller, using the same LAN that is coordinating the synchronization of the sensors. The result: data is acquired synchronously from the entire system at precise intervals, without using triggering techniques or delay compensation. There is no trigger variation caused by trigger propagation delay between the

near and far sensors. Because the 1588 modules are relatively inexpensive and can easily be added to the Ethernet-based network, the system allows for cost-effective real-time or near-real data acquisition.

PTP and Military and Aerospace Applications

Because of its precision, cost-effectiveness and ease-of-use, PTP can be expected to be widely used in military and aerospace applications. Systems that are deployed in military theater situations to identify enemy threats can be constructed from ad hoc Ethernet networks that connect a variety of sensors. Sonar systems in submarines can improve acoustic intercept and ranging accuracy by deploying PTP time stamp technology close to the sensors. And PTP is replacing IRIG systems in aircraft flight test because it provides tighter sync and eliminates the need for additional cabling.

The Future of PTP

PTP has justifiably received considerable attention since its introduction in 2002, and its influence is growing. A variety of vendors are producing hardware that supports PTP, including Intel, which has recently embedded an IEEE 1588 TSU into one of its networking microprocessors. The next version of the 1588 protocol is currently being defined and is expected to increase accuracy even more. And improved fault tolerance and compatibility with SNMP are expected to enhance PTP interoperability with existing network infrastructure. With its nanosecond accuracy, ease-of-deployment, and cost-effectiveness, PTP is poised to transform the landscape of time-synchronized applications in any number of fields.

To learn more about IEEE 1588, visit http://ieee1588.nist.gov/
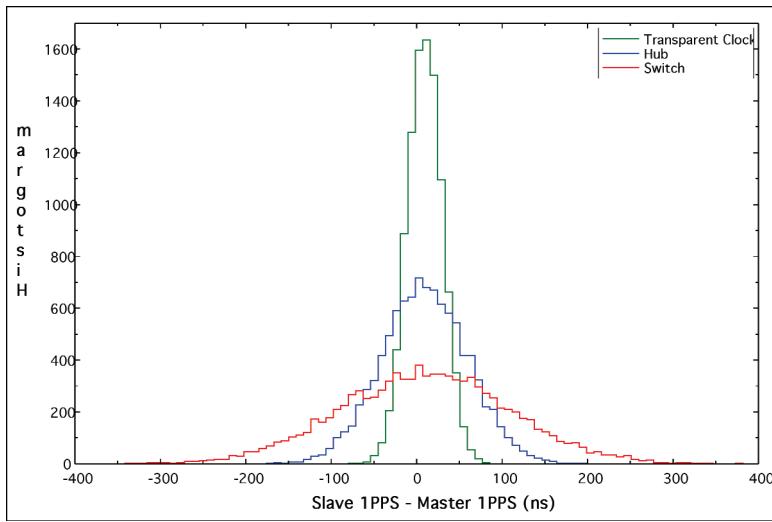
Figure 4 Histograms showing the precision of time synchronization between IEEE 1588 master and slave using three different kinds of Ethernet connection devices.

Synchronization vs. Syntonization

(Time vs Frequency)

Some applications require the time-of-day (TOD) be the same as other devices, called synchronization. Other applications require that there be a precise frequency between two or more devices, called syntonization. The means by which these are most commonly achieved in applications are as follows.

Time Transfer Mechanism	Synchronization	Syntonization
GPS	✓✓✓	✓✓✓
IEEE 1588	✓✓✓	✓✓✓
IRIG	✓✓	
NTP	✓	

✓ = Good   ✓✓ = Better   ✓✓✓ = Best