

Precision Time Protocol (PTP / IEEE-1588)

The Precision Time Protocol, as defined in the IEEE-1588 standard, provides a method to precisely synchronize computers over a Local Area Network (LAN). PTP is capable of synchronizing multiple clocks to better than 100 nanoseconds on a network specifically designed for IEEE-1588. A Network Time Server with PTP is typically referred to as an IEEE-1588 Grandmaster or PTP Grandmaster. This paper describes basic principles of PTP, information on using dual Grandmasters, and the requirements of the network to achieve sub-100 nanosecond time synchronization.

PTP Synchronization Basics

The protocol defines synchronization messages used between a Master and Slave clock similar to the Server and Client mode used in the Network Time Protocol (NTP). The Master is the provider of time, and the Slave synchronizes to the Master. A Grandmaster is a Master that is synchronized to a time reference such as GPS or CDMA.

Messages in the protocol include Master sync message, Master delay response message, and the Slave clock delay request messages. In addition to the messages, the Best Master Clock (BMC) algorithm allows multiple Masters to negotiate the best clock for the network.

Clock synchronization on the LAN requires at least one Master and one Slave. Multiple Slaves can synchronize to a single Master. The Master clock provides synchronization messages that the Slaves use to correct their local clocks. Precise timestamps are captured at the Master and Slave clocks. These timestamps are used to determine the network latency which is required to synchronize the Slave to the Master. There is a sync message transmitted typically every two seconds from the Master, and a delay request message from the Slave less frequently, about one request per minute.

Four timestamps are captured between the Master and Slave clock. The timestamps are required for the Slave offset calculation. The timestamps are commonly referred to as T1, T2, T3, and T4 (see Figure 1).




Figure 1. PTP Timestamps (T1 - T4)

Two delay paths must be calculated, the Master to Slave and the Slave to Master. First find the Master to Slave difference:

The first timestamp is T1. It is the precise time of the sync message from the Master. This timestamp is sent in the follow-up message since the time of T1 was sampled when the sync message was transmitted on the Ethernet port.

The second timestamp is T2. It is the precise time of the sync message as it is received at the Slave.

The Master to Slave difference can be calculated once T1 and T2 are available at the Slave:

$$\text{Master to Slave difference} = T2 - T1$$

Second, find the Slave to Master difference:

The third timestamp is T3. It is the precise time of the delay request message from the Slave. The fourth timestamp is T4. It is the precise time of the delay request message when received at the Master.

The Slave to Master difference can be calculated once T3 and T4 are available at the Slave.

$$\text{Slave to Master difference} = T4 - T3$$

The one-way delay can be calculated once the Master to Slave and Slave to Master difference is available at the Slave:

$$\text{One way delay} = (\text{Master to Slave difference} + \text{Slave to Master difference}) / 2$$

The offset is used to correct the Slave clock:

$$\text{Offset} = \text{Master to Slave difference} - \text{One way delay}$$

or

$$\text{Offset} = ((T2 - T1) - (T4 - T3)) / 2$$

Therefore, the following statements are true with respect to this algorithm, assuming constant network propagation delays and gradually changing operating conditions such as temperature:

The Slave clock utilizes the offset to adjust the time to agree with the Master clock. Typically, the Slave clock will use a clock tuning algorithm that can account for network propagation delays affecting the offset and the Slave clock crystal temperature and aging effect on its stability.

PTP Implementation Methods

The IEEE-1588 protocol does not define how to implement PTP into a Master or Slave. Two methods have been adopted for PTP over Ethernet: PTP software timestamps and PTP hardware timestamps. The following paragraphs describe these methods.

Grandmaster with Hardware Timestamping

While locked to GPS, the Grandmaster clock can provide precise nanosecond timestamp resolution and accuracy better than 30 nanoseconds referenced to GPS. A Grandmaster clock incorporates a local reference oscillator that is disciplined to GPS. This oscillator is the reference clock used with dedicated hardware for the precise timestamp of the incoming delay request and outgoing sync packets. Using an oscilloscope, the 1PPS (1 pulse-per-second) output from the Grandmaster can be compared to a 1PPS output from a hardware slave (see below) to measure synchronization accuracy. The dedicated hardware approach is unaffected by operating system or network traffic latency.

Slave with Hardware Timestamping

Hardware timestamps with a PTP software daemon provide precise nanosecond timestamp resolution with dedicated hardware typically in a PCIe form factor (see Figure 2). The hardware slave solution has many advantages over the software slave (see below) such as an improved oscillator, a 1PPS output for measurements compared to the master (see above), and dedicated hardware that is unaffected by the operating system latency. Synchronization of better than 100 nanoseconds is achievable using either a crossover cable or a 1588 Ethernet switch.




Figure 2. PTP Slave with Hardware Timestamping

Slave with Software Timestamping

Software-only implementations utilize existing computer hardware and a PTP daemon (see Figure 3). The slave software solution must compensate for the internal oscillator on the computer motherboard using software timestamping. The local oscillator on the motherboard is typically of poor quality and the software timestamping is affected by the operating system latency. Measuring the software slave to the master is limited to log file statistics, as there is no 1PPS output to compare with the master. Synchronization of 10 microseconds is achievable with a software slave to a master with typical results between 10 to 100 microseconds.




Figure 3. PTP Slave with Software Timestamping

Network Requirements to Achieve Sub-100 Nanosecond Synchronization

Obtaining sub-100 nanosecond timing on a local area network requires an architecture that is fully IEEE-1588 compliant. The three main components are: a GPS Grandmaster clock, an Ethernet switch (transparent or boundary clock), and a PTP slave. All components must support hardware timestamping. The Grandmaster and Slave are discussed in the [óPTP Implementationô](#) section. Below is b

Testing Synchronization Accuracy

For precise measurements the Grandmaster clock and the PTP Slave would each have a 1PPS output to measure the phase offset with an oscilloscope. See Figure 4. The synchronization accuracy and stability are dependent on the network architecture. Switches that are not 1588-compliant will add packet delay variation and impact the slave to master accuracy and stability. A crossover cable between the master and slave can be used in place of the switch to create a baseline and better understanding of the switch impact. It is also important to test synchronization capability at the intended network traffic load.

Dual-Redundant Grandmaster Clocks

Setting up the 1588-timing network with Dual-Redundant Grandmaster Clocks requires you to install and configure each of two Grandmasters according to the Best Master Clock (BMC) algorithm. See Figure 5. The IEEE 1588-2008 (v2) standard defines the BMC algorithm.

The BMC compares the clock data of the two Grandmasters and determines the preferred clock. The clock data defines the clock characteristics and is available in the announce packet. The Grandmaster with the preferred clock data becomes the PTP Grandmaster on the network and provides timing to the Slaves. Only one PTP Grandmaster provides synchronization packets at any given time. The BMC uses the clock data to determine the Best Master Clock in the following hierarchical order:

1. Priority1: User configurable from 0-255. Lower value preferred.
2. Clock Class: Dependant on the clock being locked, in holdover, unlocked, and the epoch, either PTP or ARB. In preferred order: Locked PTP epoch, Holdover PTP epoch, Locked ARB epoch, Holdover ARB epoch, Unlocked PTP epoch, Unlocked ARB epoch.
3. Clock Accuracy: Dependant on the clock accuracy (Estimated Time Error), in preferred order: <25ns, <100ns, <250ns, <1us, <2.5us, <10us, <25us, <100us, <250us, <1D □ M M T,

