



Whitepaper



IEEE 1588 PTP clock synchronization over a WAN backbone

A field study comparing PTP clock synchronization accuracy against GPS external time reference in a live production WAN environment

Contents

Background	3
PTP	3
NTP	4
GPS	4
Trial	6
Startup	7
Synchronization Quality	8
Latency	10
Routing Stability	12
Conclusions	14

Background

Endace conducted a trial of the Precision Time Protocol (PTP), also known as IEEE Std 1588™-2008, across a large corporate WAN. Clock synchronization was tested across an operational backbone. The goal of the trial was to characterize the performance of PTP as a method for providing time synchronization to sites that do not have direct access to external clock references such as GPS.

PTP

PTP is a protocol developed by the IEEE that enables the accurate synchronization of a set of peer clocks. It is intended to provide much more accurate synchronization than is normally achieved with the Network Time Protocol (NTP). PTP can provide sub-microsecond synchronization under ideal conditions.

PTP supports a number of network technologies, including Ethernet, and is primarily intended to be used on dedicated Local Area Networks (LANs). Any active network element, such as a switch or router that is added to a network, increases latency and jitter; reducing the quality of the clock synchronization that is possible. PTP aware switches, known as PTP Transparent Switches or Boundary Clocks, can reduce the impact of switching and queuing latency in the PTP network. At present, these features are not available in all mainstream switches and routers, and are typically disabled by default. Care must be taken to enable and configure PTP support in all active network elements for optimal performance.

Figure 1 shows the distribution of clock error for PTP where the GrandMaster is directly connected to the Slave. Figure 2 shows the distribution of clock error for PTP where the GrandMaster is connected to the Slave over a LAN, where the network path includes five ordinary (not transparent) Ethernet switches.

Time Source Comparison
GPS vs IEEE 1588, Direct Connection, 12 hours

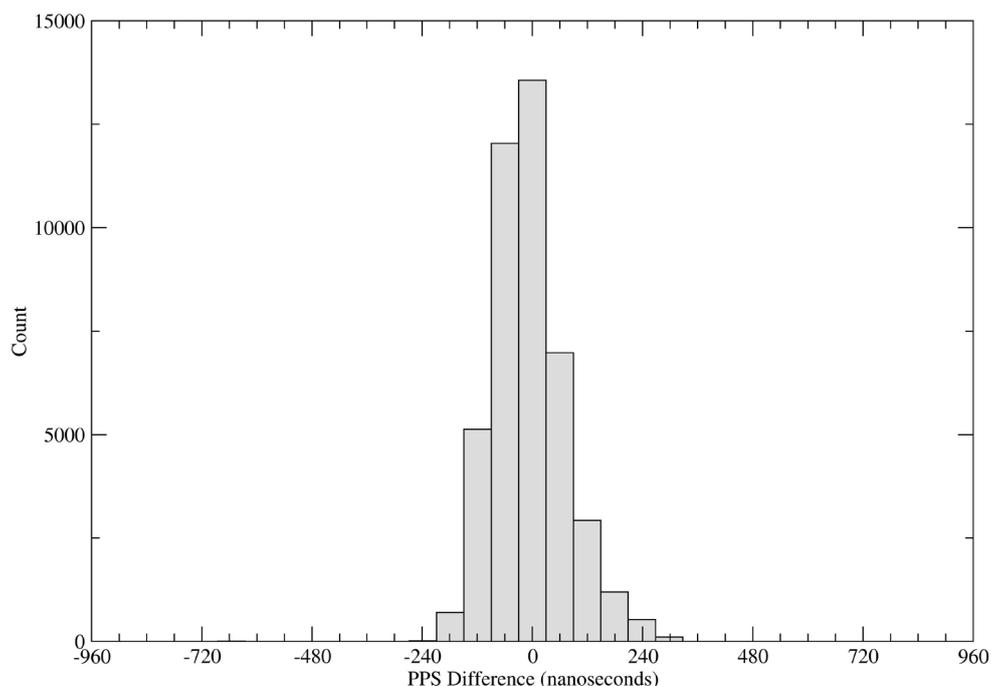


Figure 1 Measured clock error from PTP Slave directly connected to PTP Grandmaster

Time Source Comparison GPS vs IEEE 1588, 5 Switches, 12 hours

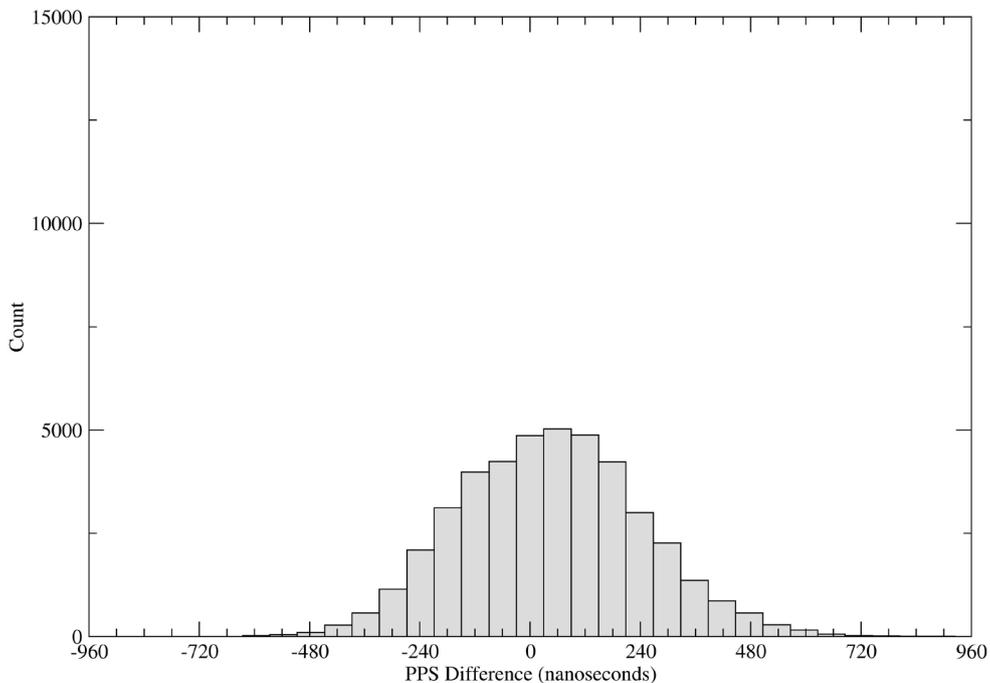


Figure 2 Measured clock error of PTP Slave connected to PTP Grandmaster via 5 switches

NTP

NTP uses a hierarchy of clocks to provide clock synchronization to about 1 millisecond over a typical Local Area Network (LAN). NTP can also be used over Wide Area Networks, including the public Internet, but with lower accuracy. Many NTP implementations synchronize the operating system (OS) clock of a PC by ‘jumping’ the clock forwards or backwards intermittently. These large changes to the clock make NTP unsuitable for accurate network jitter and latency measurements.

GPS

The US DoD NAVSTAR-GPS Global Positioning System comprises a constellation of orbiting satellites with atomic clocks on-board. As well as providing precise location information to ground receivers, it is capable of providing a very accurate time reference, typically better than 100 nanoseconds to UTC. Figure 3 shows the distribution of clock error between two Endace DAG cards synchronized to separate GPS receivers over a 24 hour period.

Time Source Stability
GPS x2, DAG 4.5G2 x2, 2 million packets, 24 hours

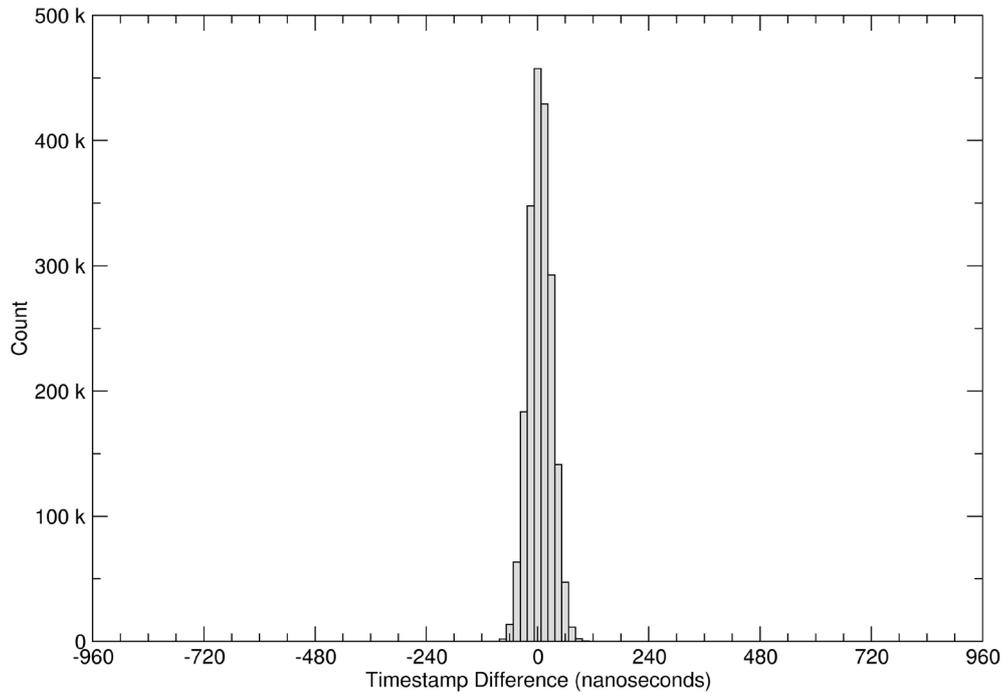


Figure 3 Comparison of two independent GPS receivers

A disadvantage of GPS for distributing time signals is that the antenna must be placed on the roof with a clear view of the sky in order to reliably receive the satellite signals. It can be expensive, time-consuming, and logistically difficult to get a GPS receiver mounted on the roof and cabling run to the appropriate floor to the receiver. At some shared sites it may not be possible to get approval at all.

Trial

For this trial, Endace was testing the feasibility of using PTP to provide clock synchronization at sites with no GPS receivers, by sending synchronization messages across the WAN network from sites that do have GPS receivers.

In New York, two sites that currently have GPS receivers installed were identified, site A and site B. By sending PTP messages between these sites over the WAN network it was possible to compare the performance of PTP at the slave site to the local GPS reference.

Endace supplied a 1RU system to operate as a PTP master at site A. An existing Symmetricom XLi GPS receiver was used as the clock source for the PTP Master. A second Endace 1RU system was installed at site B to act as the PTP Slave. The Slave generates a 1 pulse per second (1PPS) timing signal and this was compared to the Symmetricom XLi GPS receiver to determine the clock synchronization accuracy.

A full 24 hour period from 00:00 to 23:59 EDT was chosen for analysis. By comparing times during the day and evening it was possible to determine the effect of operational traffic on the clock synchronization, as the same physical network is used for both. The WAN network is OSPF routed and fully redundant. The primary path between sites A and B includes two Cisco 3750 Gigabit Ethernet switches at the end points, and four Juniper M320 routers connected by 10 Gigabit Ethernet.

Startup

First, the PTP Master at A synchronizes its clock to the local Symmetricom XLI using a hardware 1PPS signal. This takes less than two minutes. After initial synchronization, logs showed that the PTP Master remains synchronized to within ± 45 nanoseconds of the Symmetricom during the entire trial period.

When the network routing between the PTP Master at A and the PTP Slave at B comes up, the Slave begins synchronizing to the Master. Figure 4 shows the PTP clock offset error estimate for the first 10 minutes after routing comes up. The clock error estimate (Offset from Master) starts at zero as no information about the error is available yet. The estimate is refined over the first few seconds before clock adjustment begins.

The clock error between the PTP Slave and Master was <100 microseconds within 2 minutes, <10 microseconds within 6 minutes, and <2 microseconds within 10 minutes.

Offset from Master
A to B

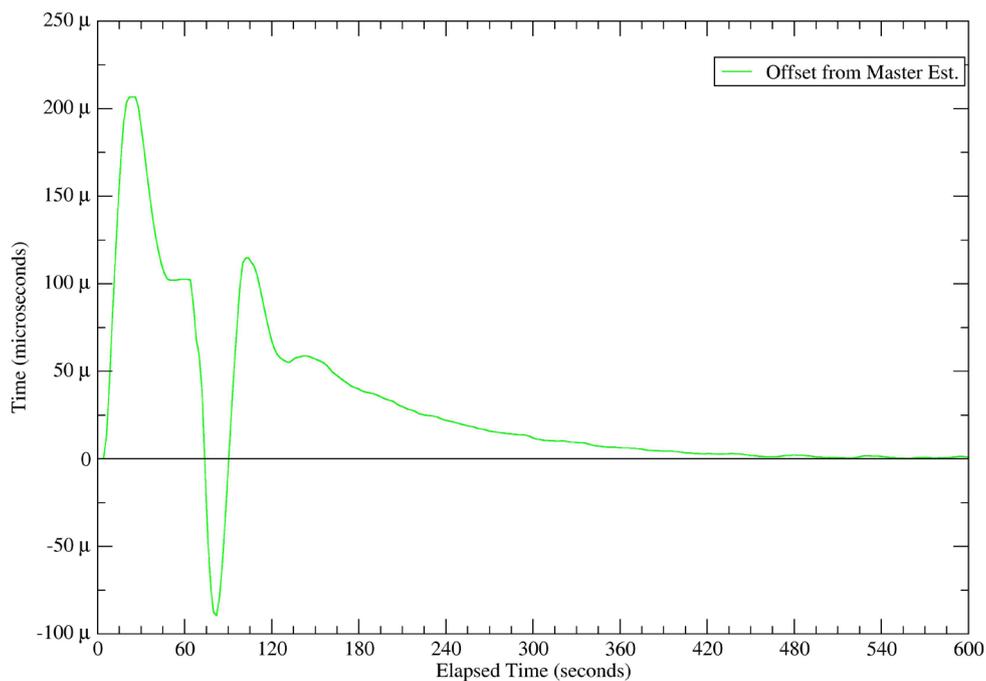


Figure 4 PTP Slave Offset from Master estimate

Adjusting PTP protocol parameters such as the update message rate can affect the initial synchronization time. For this trial, the default 2 second update rate was used.

Synchronization Quality

The quality of the clock synchronization at the PTP Slave is measured by time stamping the PTP Slave 1PPS output every second using the Symmetricom GPS in Event Timing mode, and calculating the difference between the two clocks.

Figure 5 shows the distribution of clock synchronization error at the PTP Slave over the 24 hour period. The error is centered about 0, with a standard deviation of 1.07 microseconds ($3\sigma = 3.2\mu\text{s}$). The largest outlier is almost 50 microseconds.

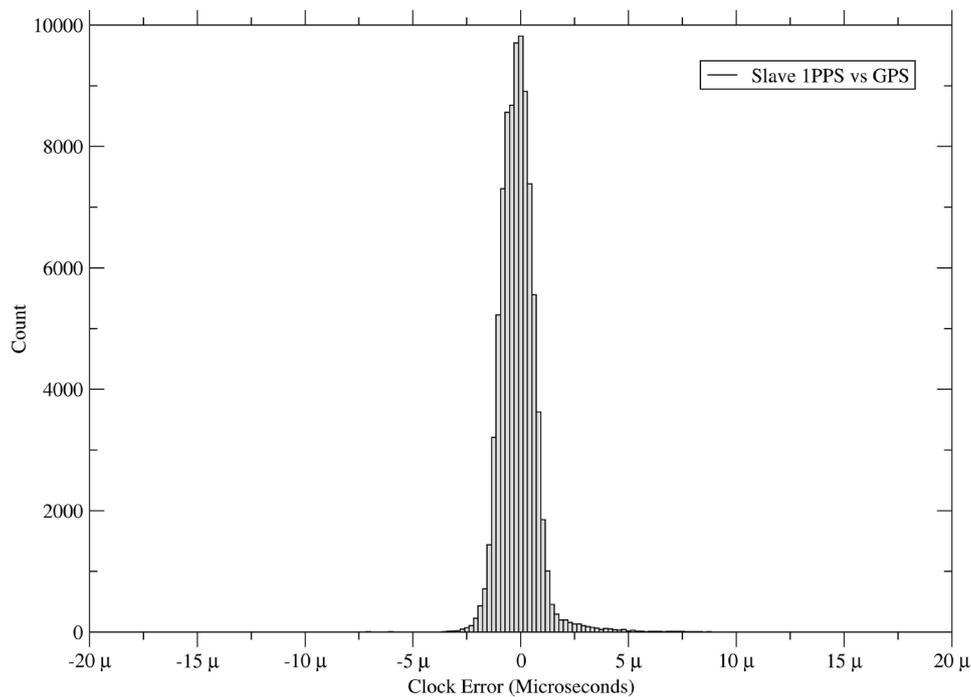


Figure 5 PTP Slave Clock offset vs GPS histogram

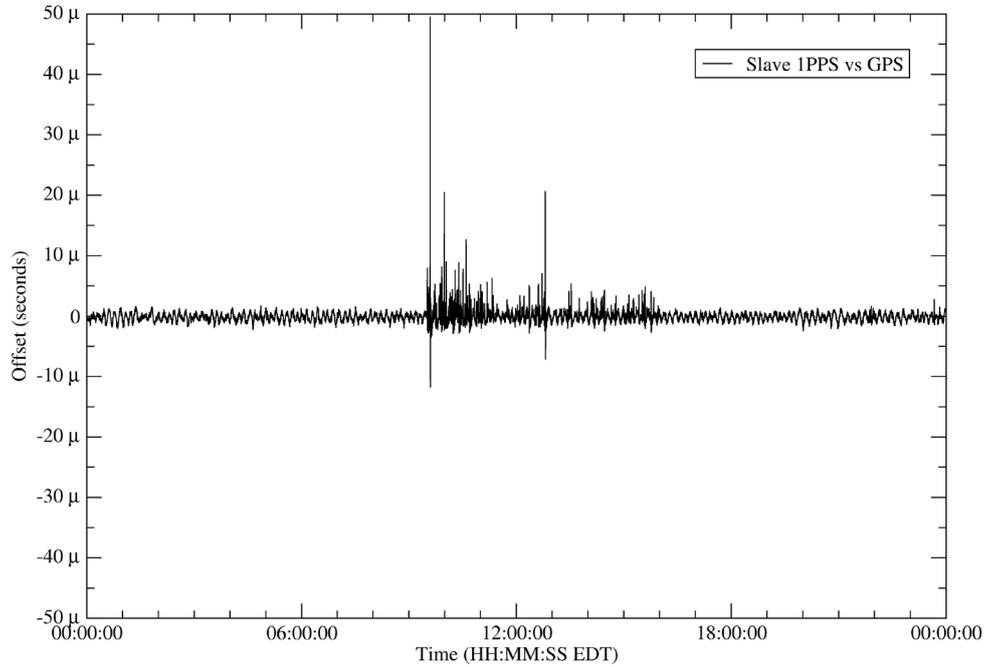


Figure 6 PTP Slave clock offset versus GPS reference

The impact of operational traffic during the day is apparent from 09:30 to 16:00 where the peak clock error is higher. The largest outlier is almost 50 microseconds. However, the mean error does not change greatly during the day. The PTP client uses non-linear filtering which has removed the majority of the clock offset outliers. Further tuning of the filter parameters to increase their depth would improve their ability to reject the outliers, and should eliminate the remaining peaks.

While operational traffic is light during the morning and evening, the clock synchronization error is typically ± 2 microseconds. SNMP logs of traffic volume on the network path showed peaks at only 300Mbps during the day, however micro-bursts at finer time scales may have been much higher.

Latency

PTP synchronizes Slave clocks to a Master by exchanging messages across a network. Packets take a finite time, or latency, to cross the network due to physical propagation and forwarding delays.

The PTP system attempts to measure the network latency in order to compensate for this delay. It accomplishes this by exchanging delay request messages across the network and comparing the time stamps for the time the packets were sent and the time they were received.

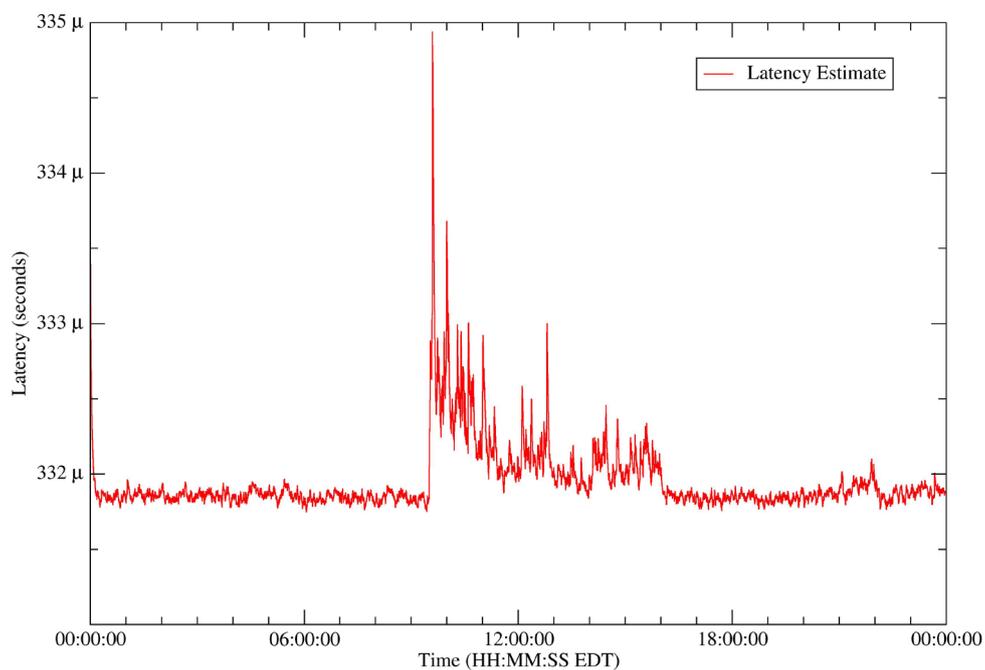


Figure 7 PTP One-way Delay estimate

Figure 7 shows the estimated network latency between A and B. This is the One-way Delay, which is computed as half of the Round Trip Time. The network latency is assumed by PTP to be symmetric. The estimated latency is very stable over the 24 hour period, at just under 332 microseconds. This latency is equivalent to the propagation of light through 41 miles (67km) of fiber. The actual fiber path is probably shorter since some of the delay is attributable to network switches and routers. The line of sight distance between sites A and B is actually less than two miles, indicating the network path is far from direct.

The latency estimate does increase during the day. This could be due to congestion at switches or routers increasing the latency, or to some packets following different network paths. It may be possible to prevent the latter with traffic engineering.

The network latency estimate is constructed by measuring the latency from the PTP Master to the Slave, and the return latency from the Slave to the Master.

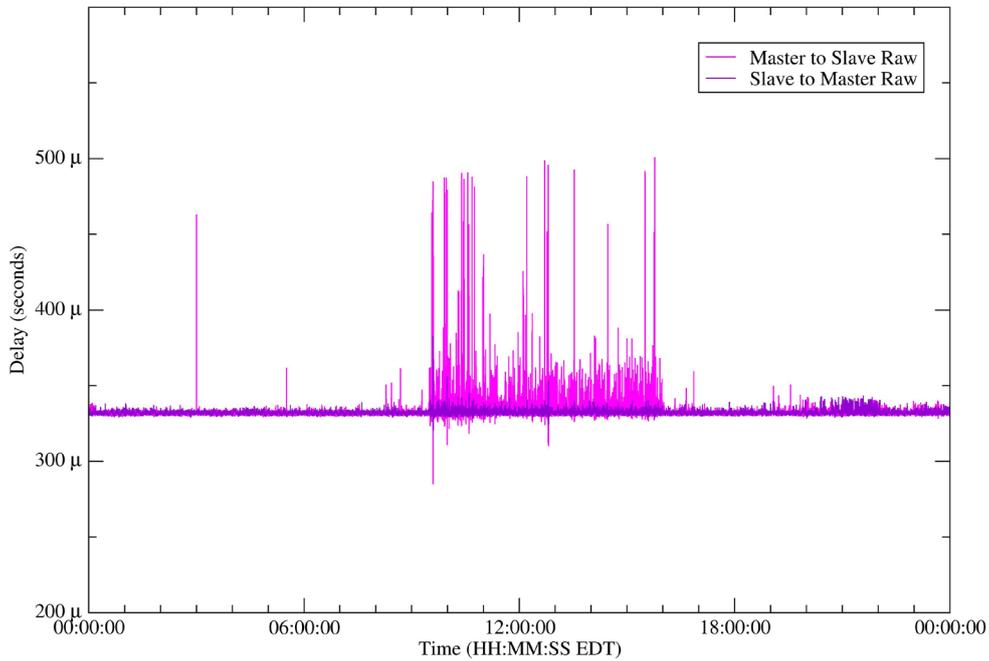


Figure 8 Raw PTP One-way Delay measurements

Figure 8 shows the raw latency measurements that were taken over the day. The raw latency measurements are passed through a median filter to reject outliers. The filtered latency measurements are shown in Figure 9.

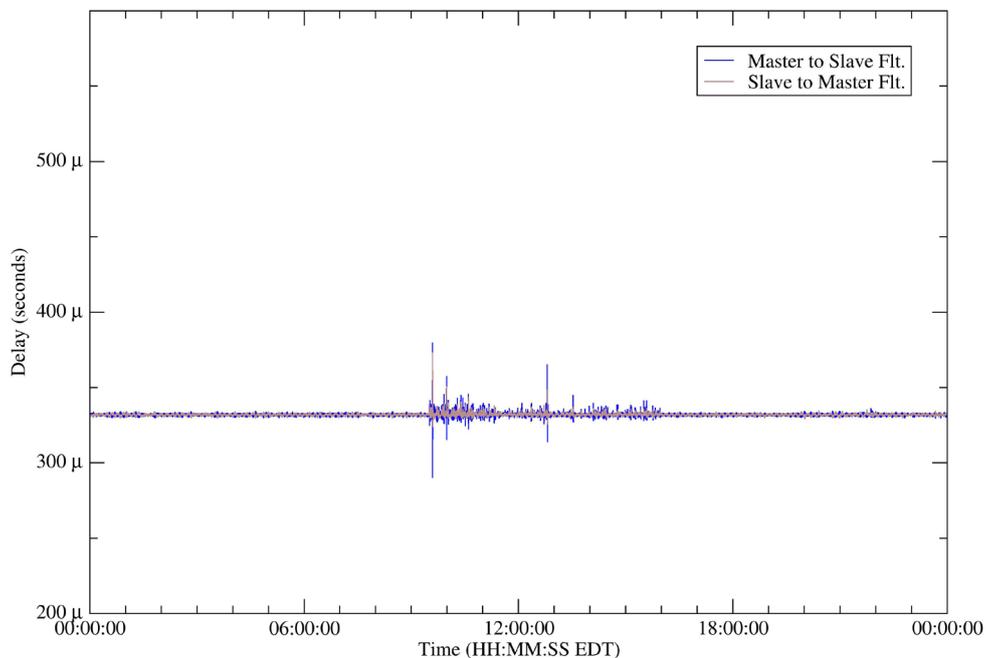


Figure 9 Filtered PTP One-way Delay measurements

The reduction in noise can clearly be seen, but there are still a few outliers. Adjusting the filter parameters should reduce the outliers further, improving the network latency estimate and the quality of the clock synchronization as a consequence.

The final network latency estimate is the result of combining and further filtering the Master to Slave and Slave to Master one-way latency measurements.

Routing Stability

In order to calculate the network latency, PTP makes a critical assumption: that the network delay is symmetric.

This assumption is reasonable for switched networks, such as Ethernet LANs, as packets typically follow the same path through the network in both directions. However, in routed networks packets can take different paths in each direction. If packets travelling in one direction always take a fixed route but packets in the reverse direction take different routes, the latency they encounter will be different and the network delay asymmetry will change.

PTP cannot distinguish changes in delay asymmetry from changes in clock offset between the Master and Slave.

This is illustrated by an event that occurred outside the selected day. At approximately 23:43 the network latency between A and B changes, as shown in Figure 10.

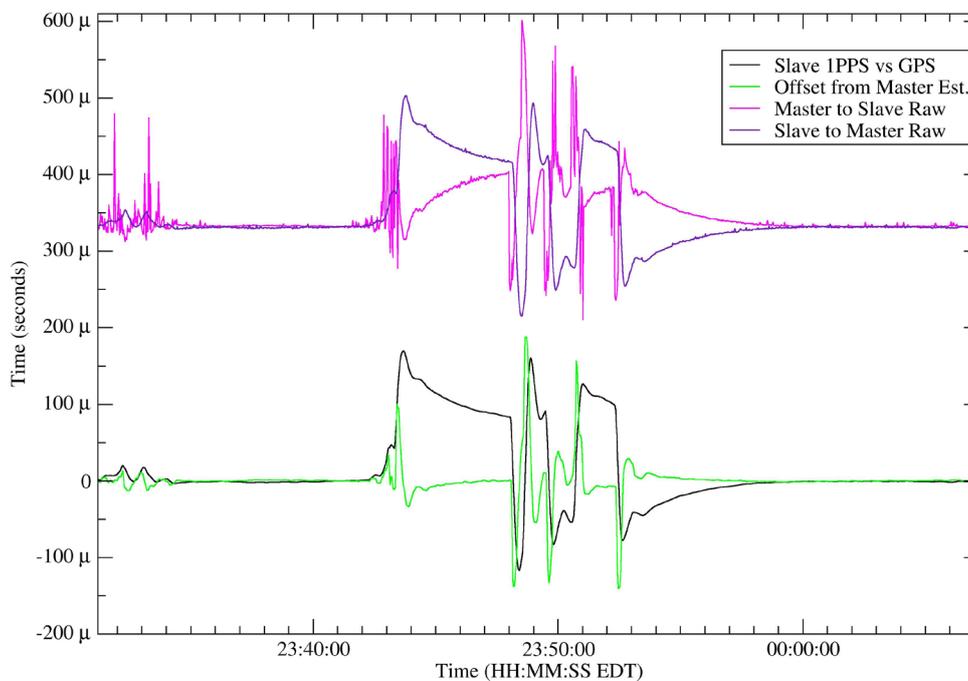


Figure 10 Raw measurements of PTP anomaly

The raw latency measurements from the Master to Slave (magenta) start to fluctuate, suddenly increasing from 332 microseconds to 500 microseconds: an increase of 168 microseconds. The PTP Slave interprets this as a slave clock offset, and rapidly skews the clock to reduce the perceived error. This clock change then affects the Slave to Master latency measurements (violet), and subsequent Master to Slave measurements. The change affects all the packets travelling on this network path and remains stable for 5 minutes. 168 microseconds is equivalent to an additional 20.7 miles (33km) of fiber propagation.

The change in network latency affects the clock synchronization, causing an error of 168 microseconds (pps, black). Because the latency of the path changes in only one direction, the delay asymmetry also changes. PTP cannot distinguish the change in asymmetry from a clock offset error. After five minutes the clock error is equal to one half of the delay asymmetry, or 84 microseconds.

The Master to Slave latency then reverts to its original value for a few seconds, causing the clock offset error to reverse. Over the next five minutes the network latency changes four more times, and eventually returns to its original symmetric value. The remaining clock offset error is then averaged down to zero before 00:00.

In Figure 11, we compensate for the changes in the Slave clock in order to show what the delay measurements would have been without the consequent clock adjustments. This makes it clearer that only the Master to Slave path delay has actually changed.

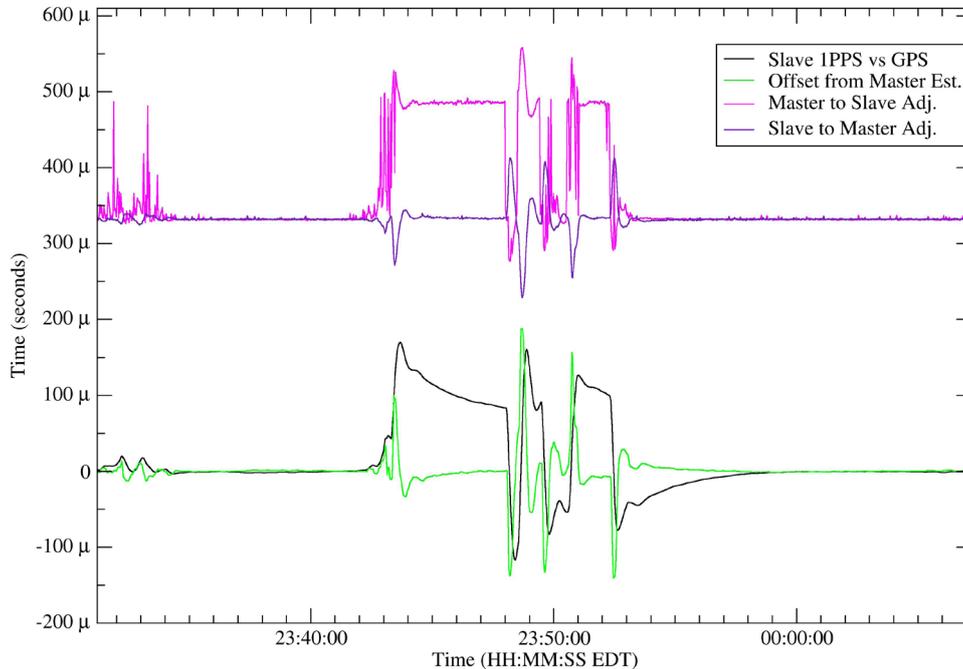


Figure 11 Figure PTP Anomaly compensated for Slave clock adjustments

There are two potential explanations for this asymmetric increase in latency: network congestion affecting one direction of the path causing queuing, or an asymmetric route change.

Cross traffic within the network could cause output port contention at one of the routers or switches when the volume of traffic routed to a port exceeds the link bandwidth. This can cause buffering within the switch, delaying packets. To account for a 168 microsecond delay, the PTP packets would have to be buffered behind 205kB of packets on a 10Gbps port, or 20kB on a 1Gbps port. In this situation we would also expect to see some packet loss (non-delivered PTP frames), but this was not observed.

An alternative explanation is a change in the physical path taken by packets from site B to site A. This might be caused by IP routing flaps, or by changes in layer 1 or 2 paths such as MPLS or SONET/SDH ring automatic protection switching.

When traffic engineering is used to ensure that only symmetric network paths are used, a link failure will cause a change to a new symmetric path. This new symmetric path will likely have a different physical path and hence a different propagation time and latency. In this situation PTP will initially show a non-zero clock offset, however it will re-converge to having zero clock offset. The re-convergence time depends on the magnitude of the network latency change, but will generally be under 10 minutes.

Conclusions

PTP is a protocol for synchronizing clocks over IP networks, which this field study has demonstrated to be viable over an operational WAN network for clock synchronization within ± 3.2 microseconds at 3σ . The largest clock error during the day was 50 microseconds, however further tuning of the filtering algorithms could reduce such outliers.

PTP is potentially a solution for providing clock synchronization services to sites that do not have direct access to external time references such as GPS, depending on synchronization accuracy requirements. PTP can be provisioned over the existing operational network, providing high quality clock synchronization without affecting the operational network. PTP is simpler to provision to a site than GPS, and may be significantly cheaper than negotiating roof access and rental.

We have seen that while PTP clients provide an estimate of the clock Offset from Master, this estimate can be significantly incorrect for extended periods under anomalous network conditions. The PTP client Offset from Master estimate cannot be relied upon to accurately indicate clock synchronization quality, or provide traceability to UTC.

Best practices for PTP deployment should be followed to reduce the impact of queueing delays. PTP systems should be calibrated for systematic errors, and closely monitored for potential anomalies. Because PTP cannot distinguish between Slave clock offsets and asymmetric changes in network latency, direct reference to a primary clock source such as a GPS disciplined clock is recommended for applications where accuracy and traceability to UTC are paramount.

Table 1 Clock distribution performance

	Best Accuracy	Worst Accuracy	Deterministic	Audit Trail
GPS	~50ns	~100ns	Yes	Yes
PTP (LAN)	< 1 μ s	> 10 μ s	Maybe?	Maybe?
PTP (WAN)	< 10 μ s	> 100 μ s	No	No