

2023-May-31



PCIe PTM: PRECISION TIME IN THE LAST CENTIMETERS

Securities Technology Analysis Center (STAC) Summit / New York City

Kevin B. Stanton, Ph.D.
Sr. Principal Engineer
Intel Corporation

DISCLAIMER

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more on the [Performance Index site](#).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

PRECISION TIME IS A FOUNDATIONAL CAPABILITY

Time Sync required
and Deadlines
must be met



Automotive



Industrial / CPS



Conferencing




Live A/V



AR/VR
and Gaming



Mobile
Robots



Time Sync is
required
[Sometimes
Deadlines too]



5G / Comms



Financial



Cloud



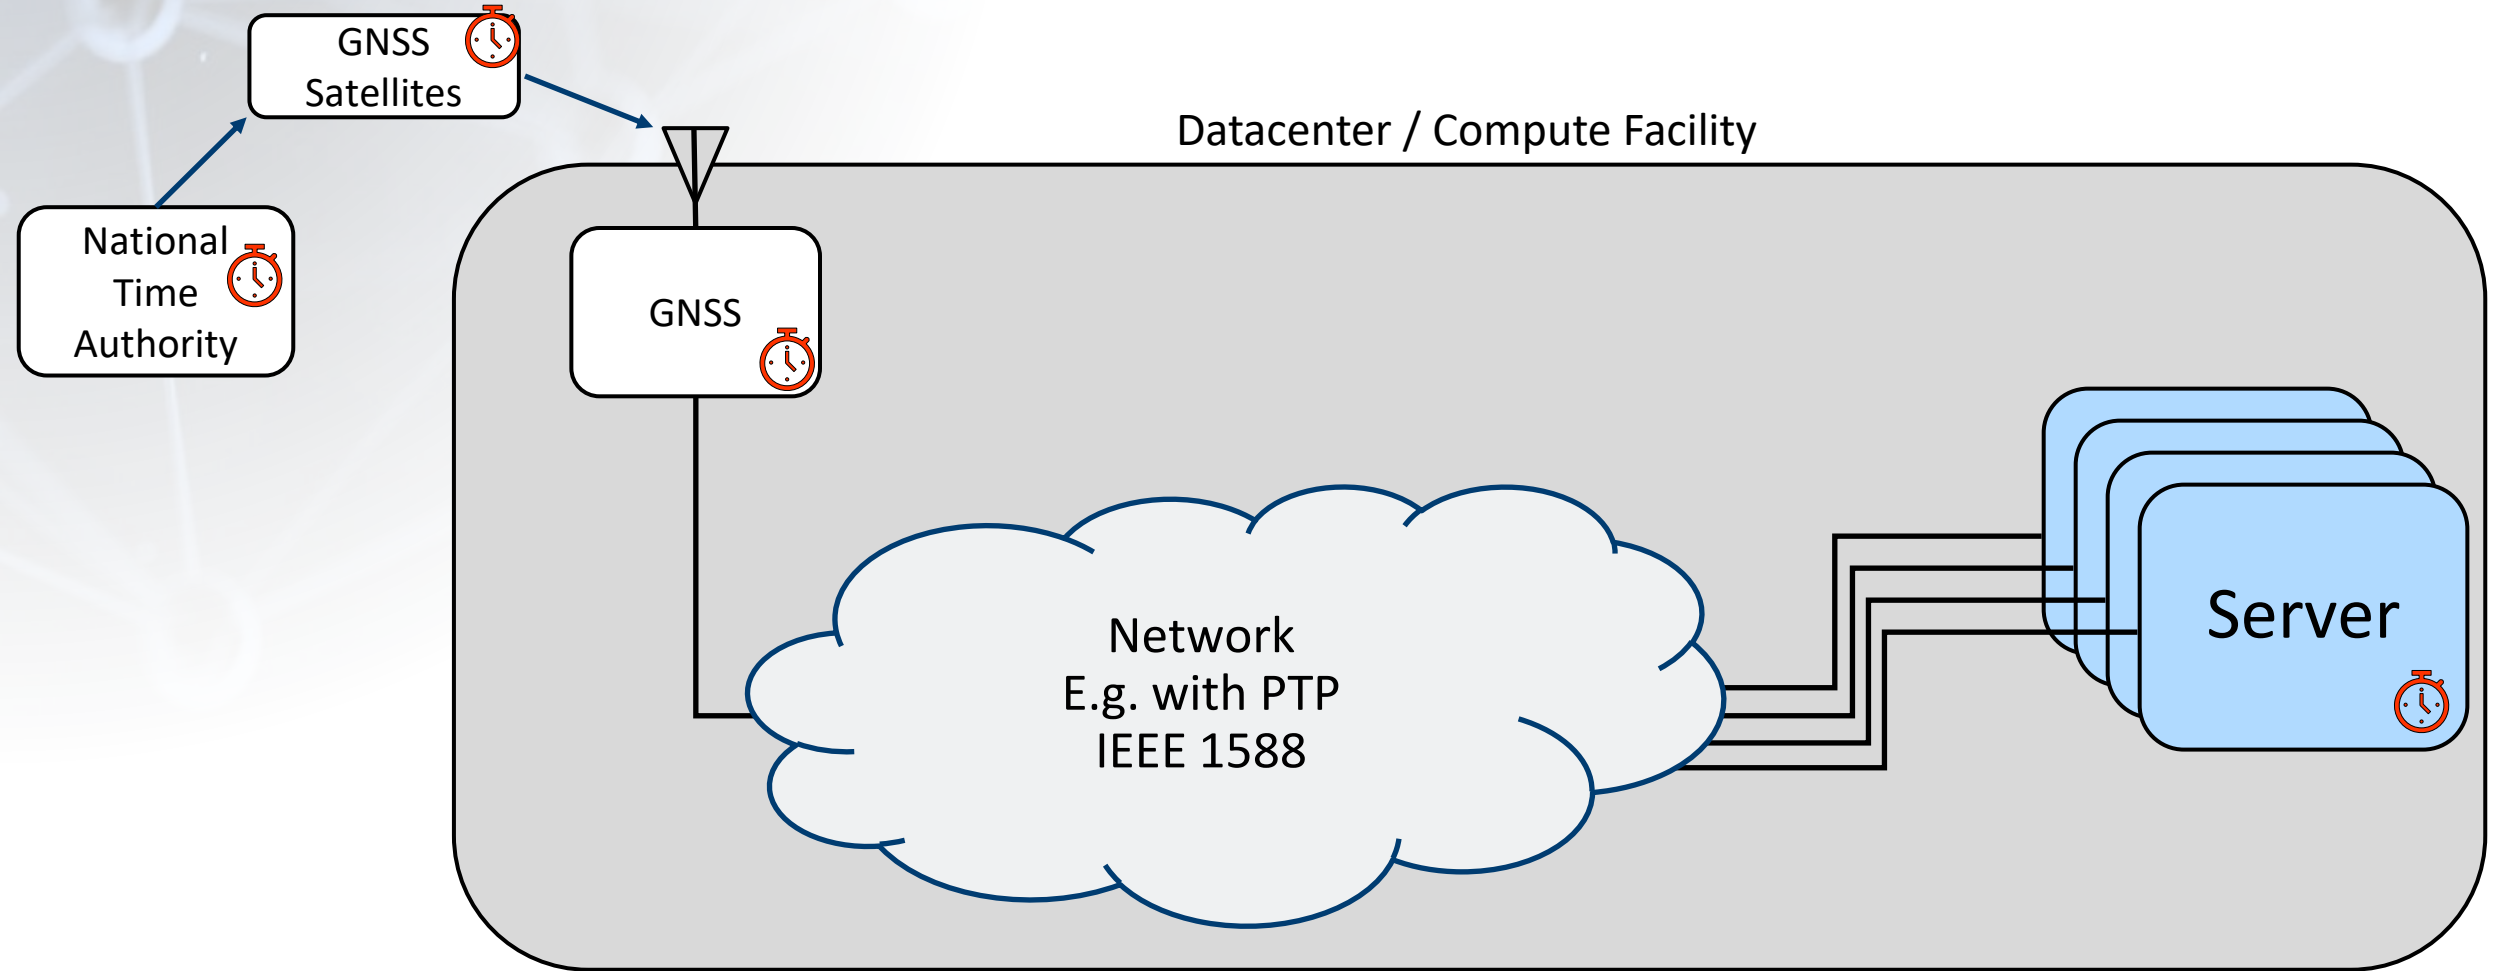
Edge



HPC

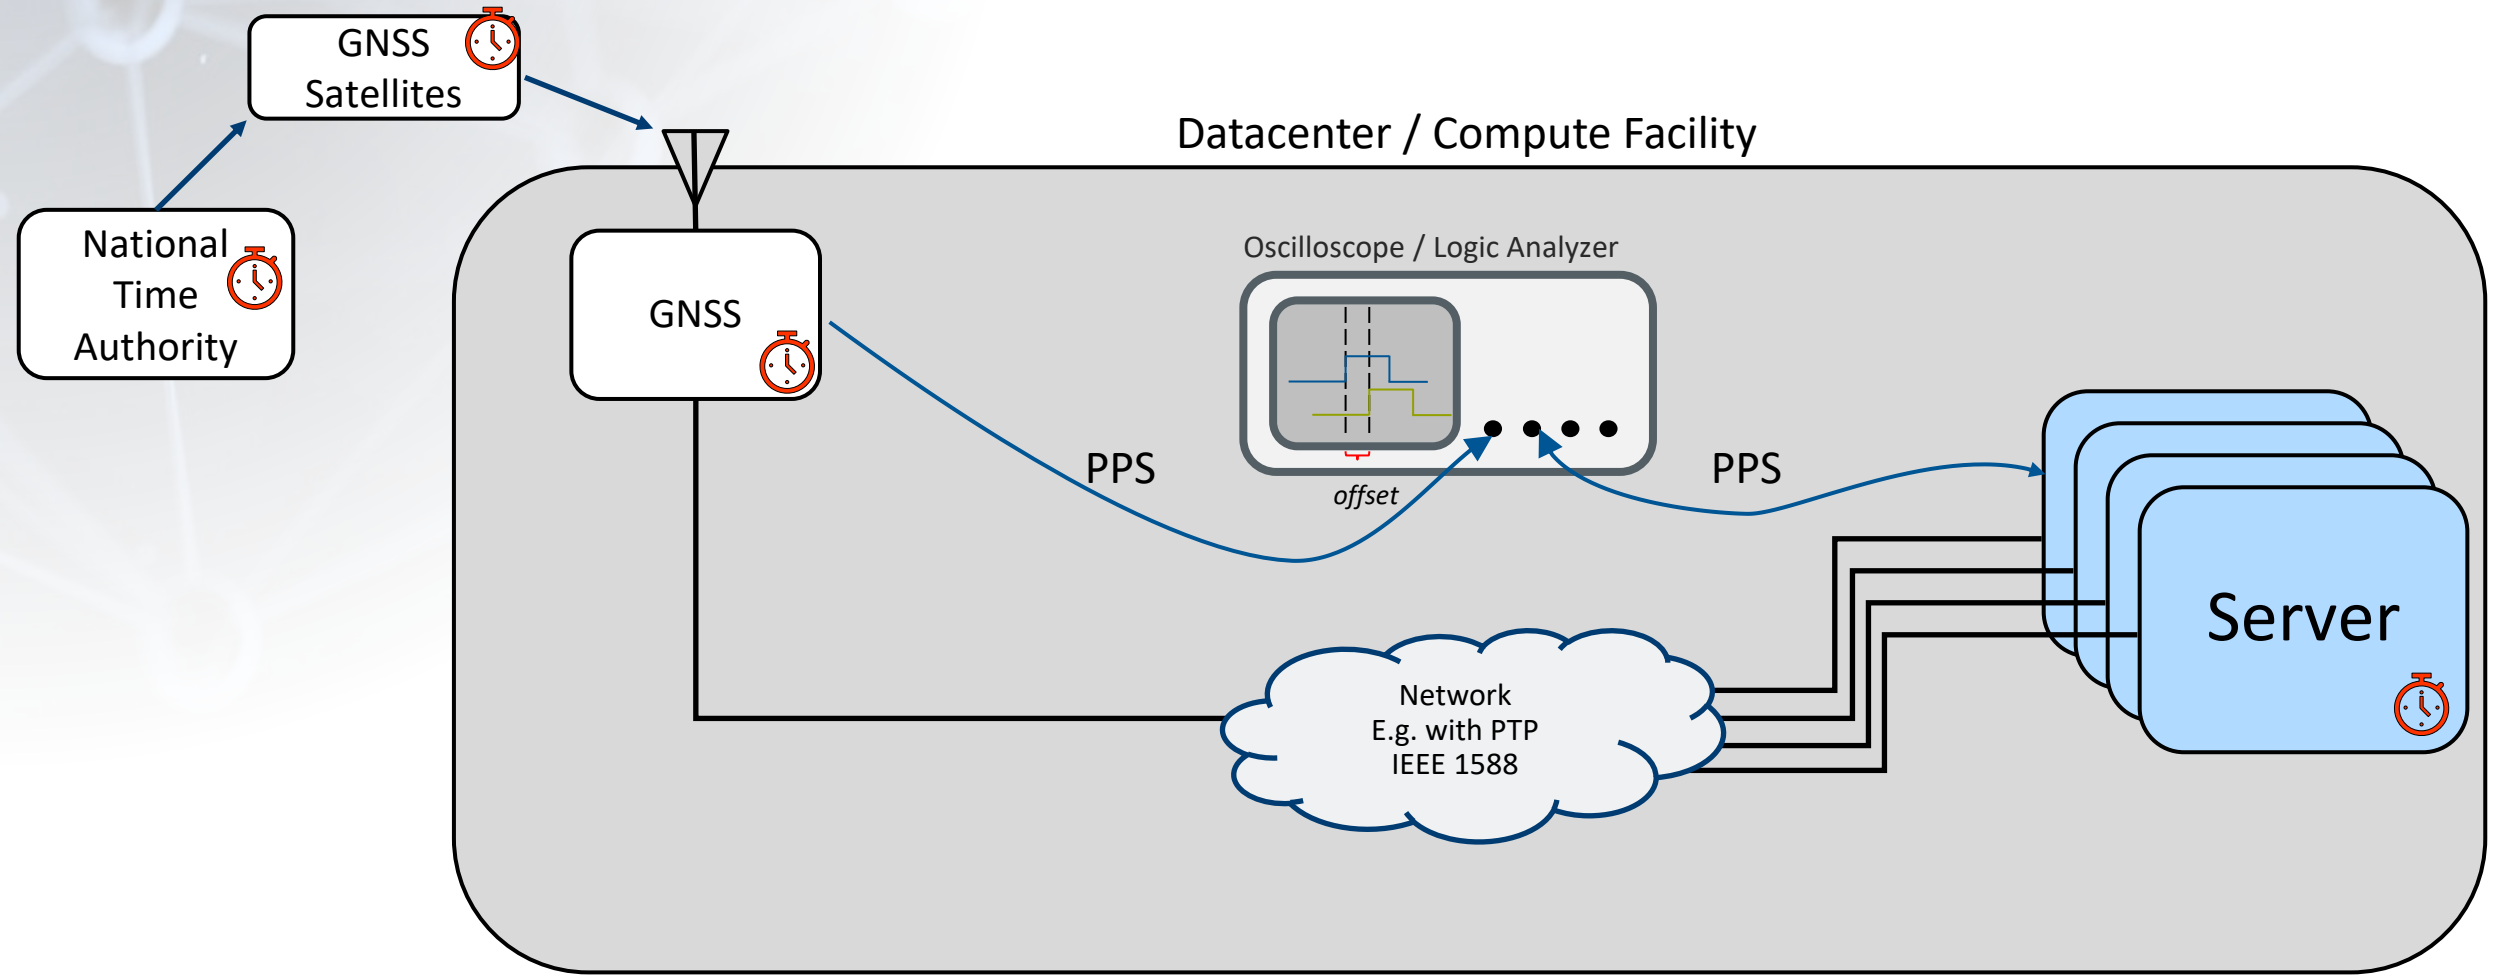
Where Time Bounds are Necessary for Correctness of the Application, Compliance, etc

TIME SYNCHRONIZED ACROSS 45,000 KM*



* Not drawn to scale

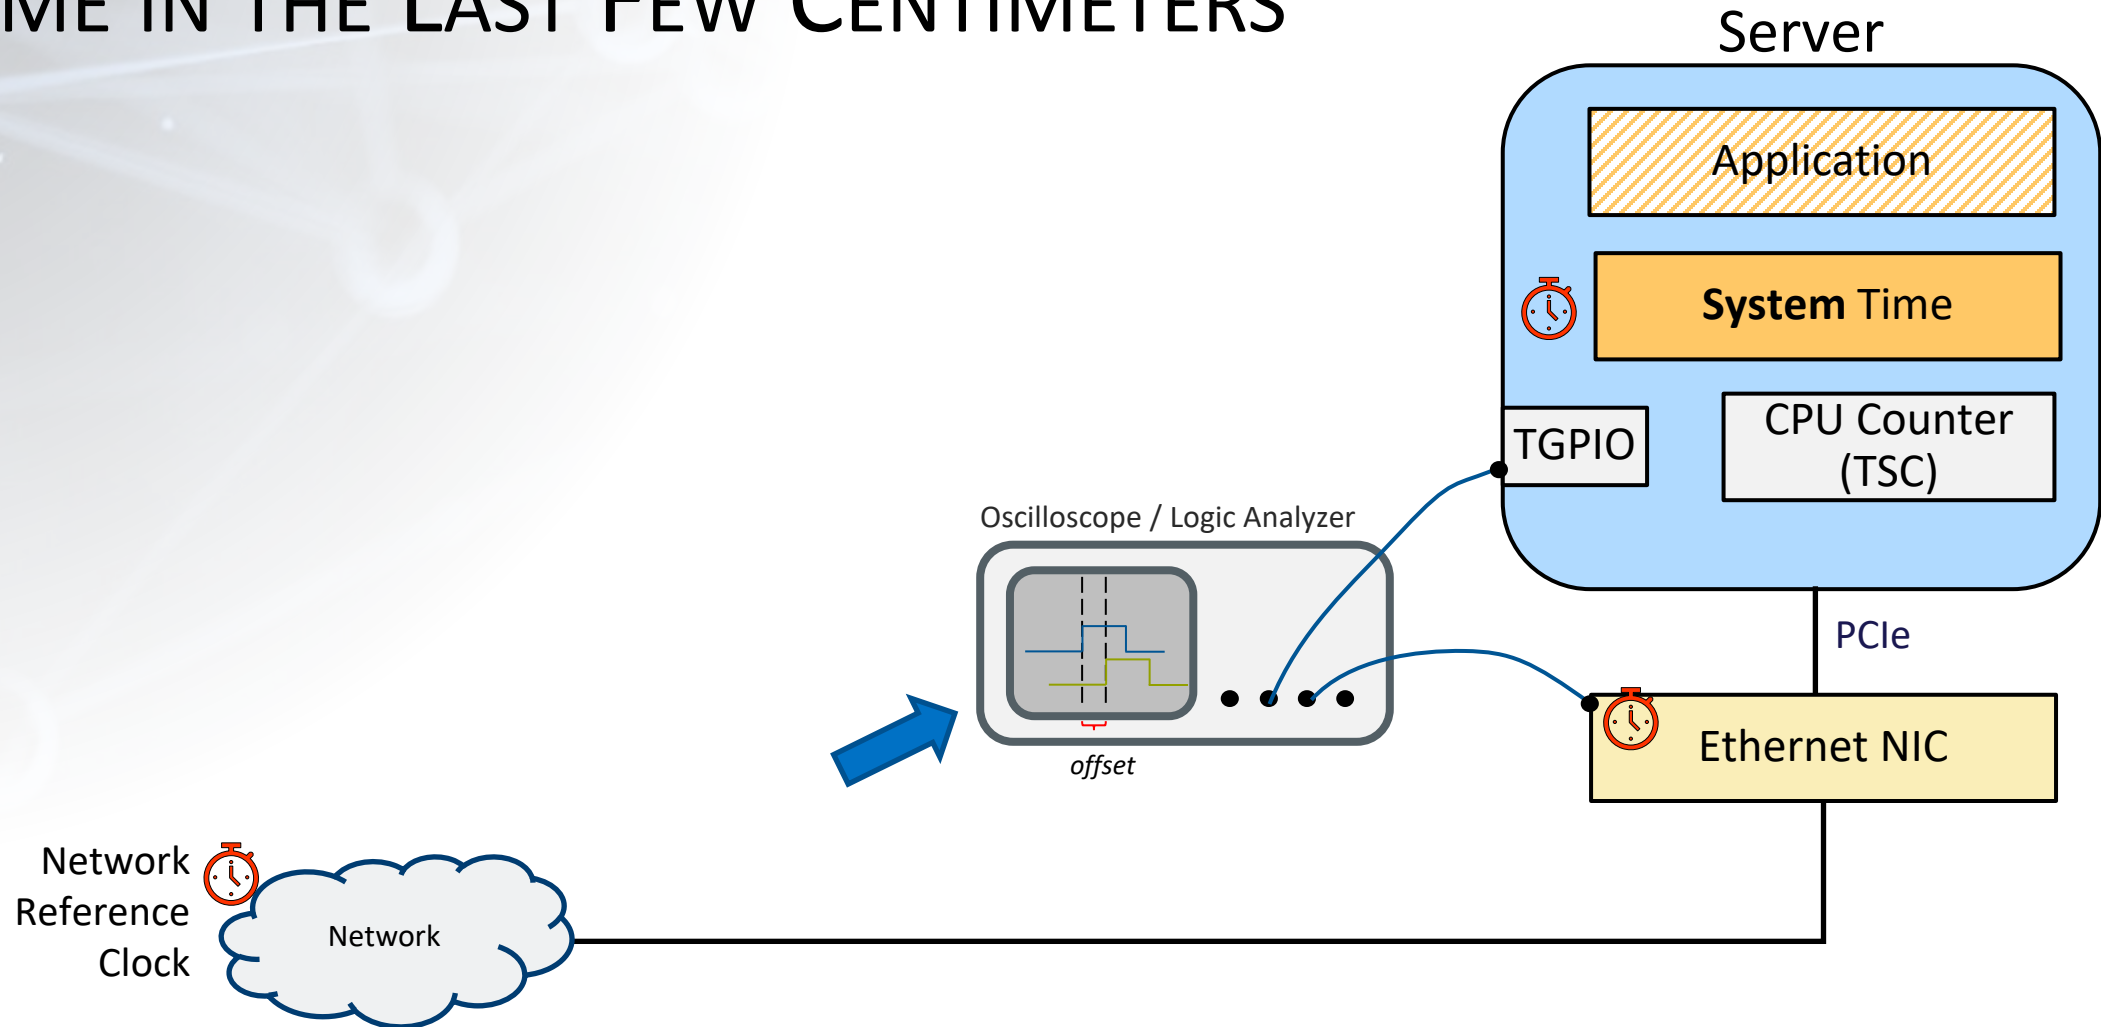
PPS USED TO VERIFY TIME SYNC ACCURACY



* Not drawn to scale

PPS = Pulse Per Second

TIME IN THE LAST FEW CENTIMETERS



TIME SYNC QUALITY IN THE LAST CENTIMETERS


System Load

Idle

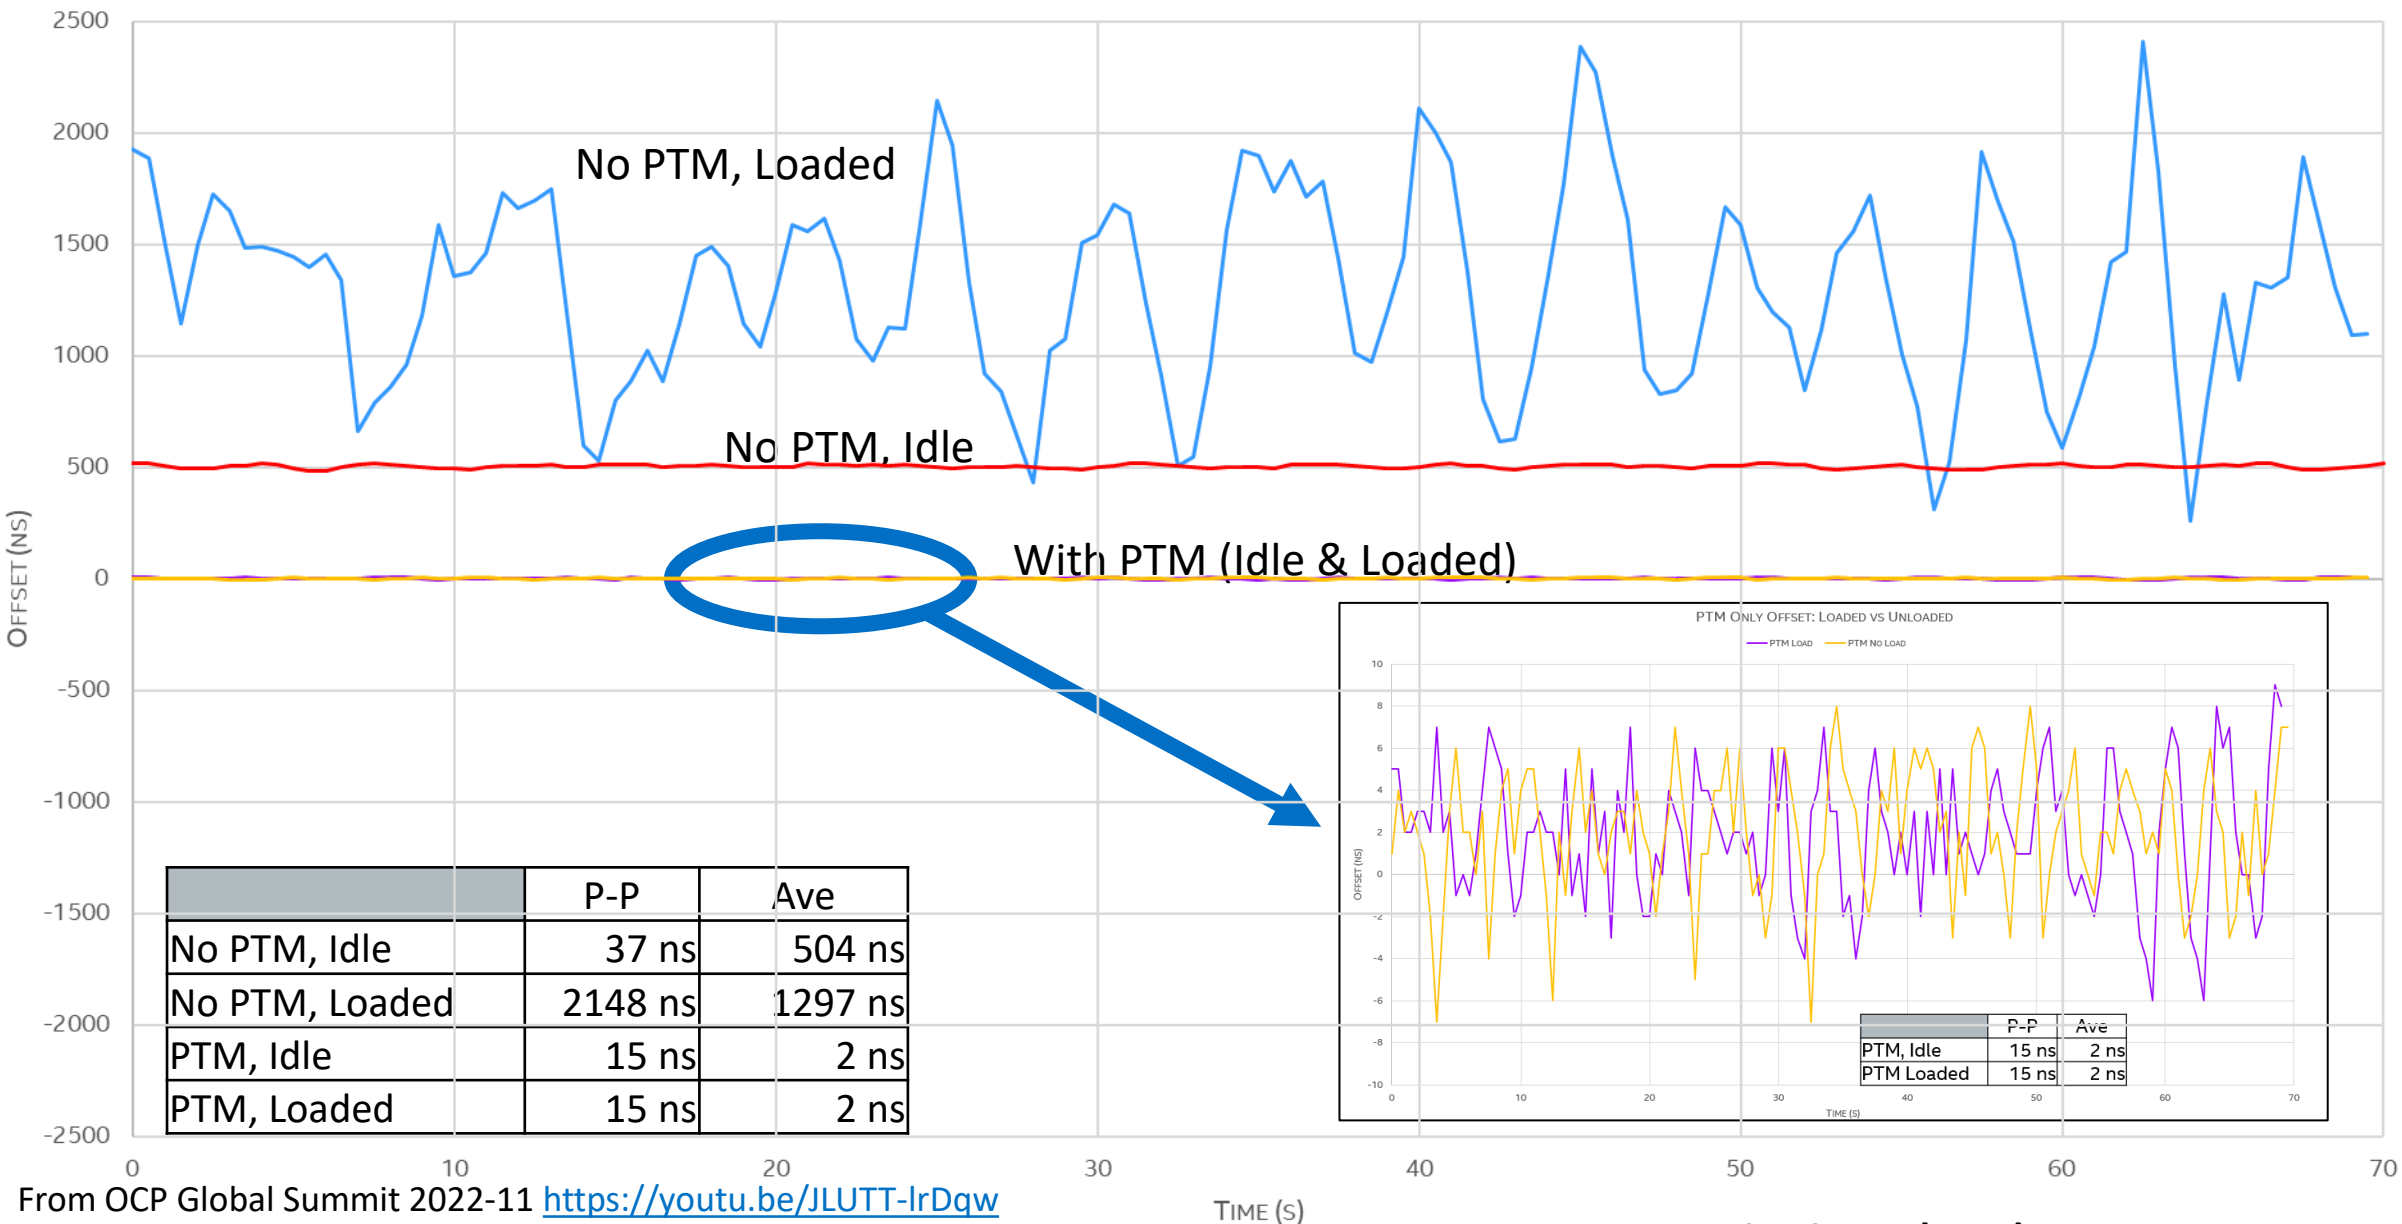
Heavily Loaded

No PTM

PTM



TIME INACCURACY WITH AND WITHOUT PCIe PTM

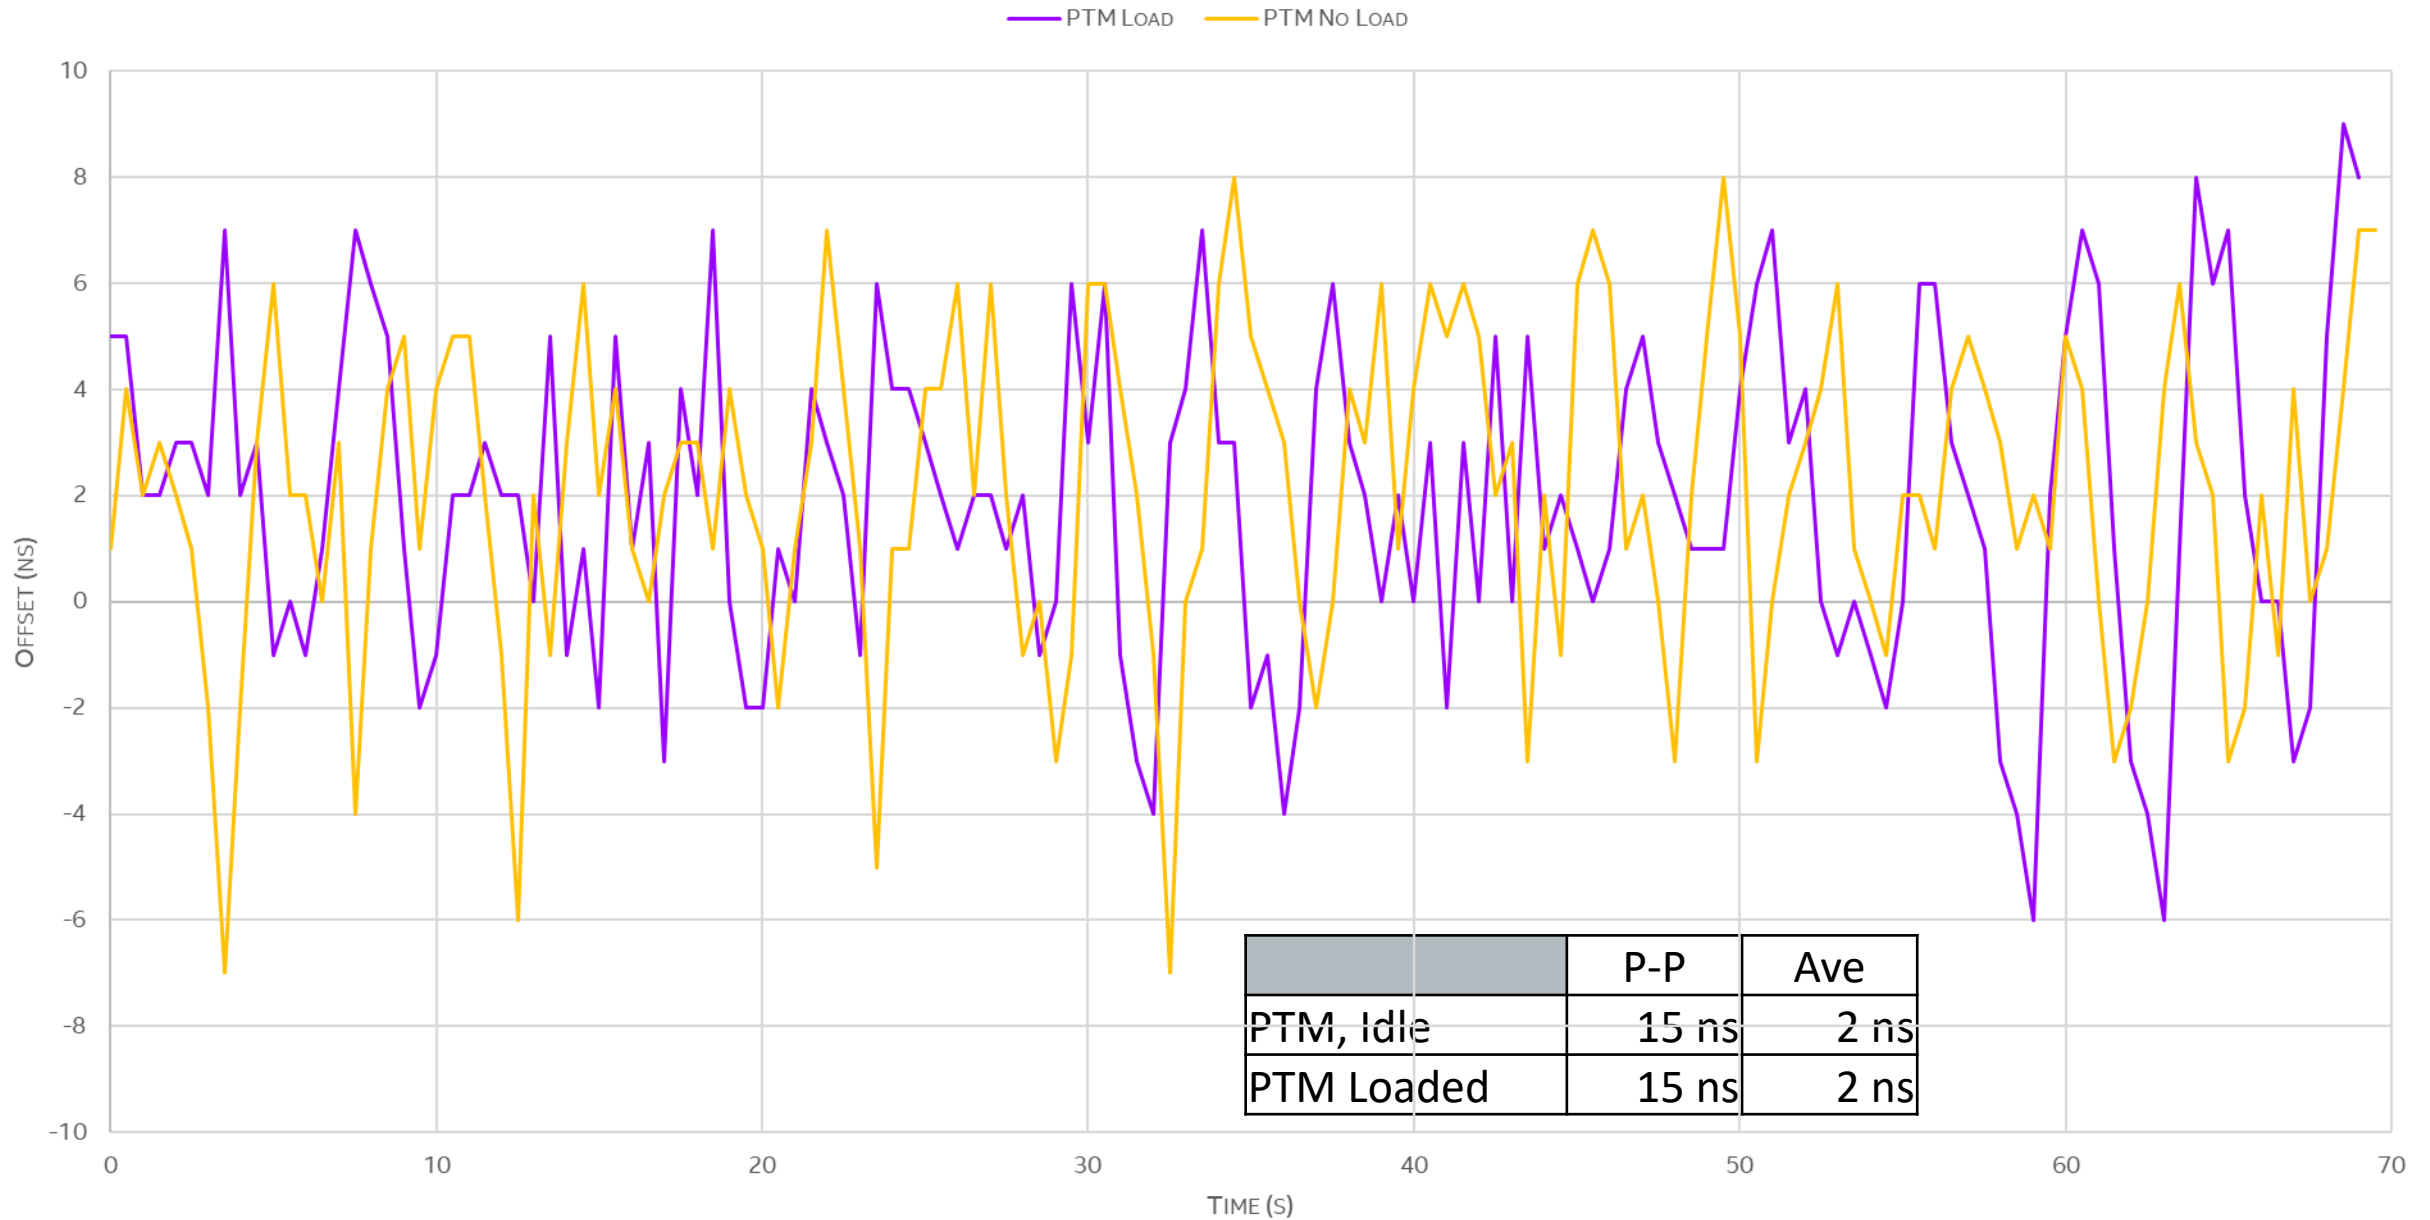


From OCP Global Summit 2022-11 <https://youtu.be/JLUTT-lrDqw>
See backup for system configuration. Results may vary.

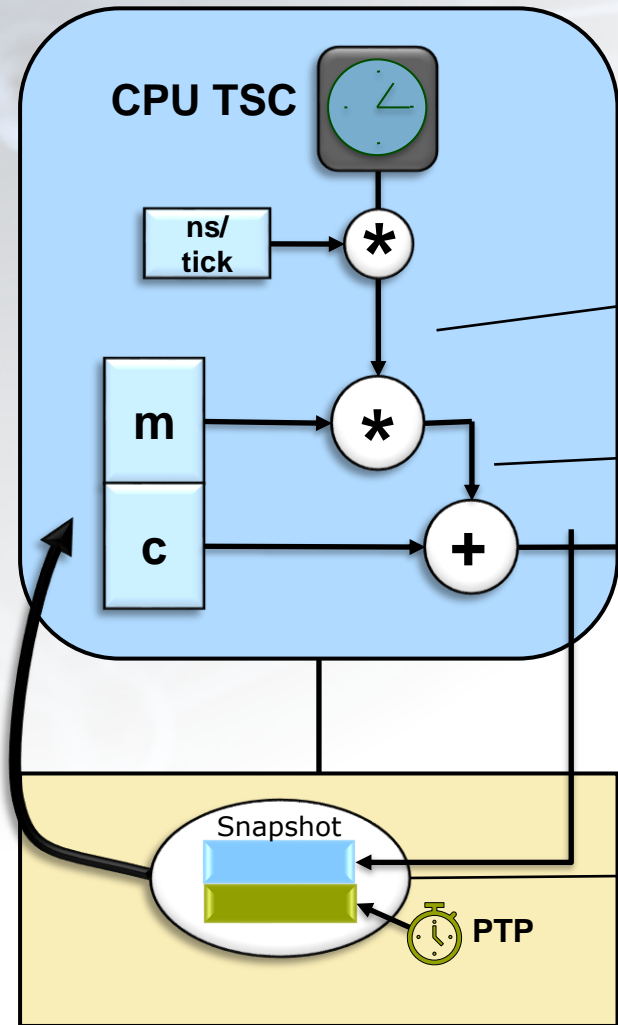
Not STAC Benchmarks

TIME INACCURACY WITH PCIe PTM (IDLE AND LOADED)

Not STAC Benchmarks



CPU COUNTER → SYNCHRONIZED TIME



Time “now” (from a Linux application)

`clock_gettime(CLOCK_MONOTONIC_RAW, &now);`

- Returns TSC scaled to nominal nanoseconds

`clock_gettime(CLOCK_MONOTONIC, &now);`

- Returns TSC scaled to track TAI, in nanoseconds

`clock_gettime(CLOCK_REALTIME, &now);`


- Returns `CLOCK_MONOTONIC + (now-1/1/1970)` [incl. leap seconds]

Cross-Timestamp “Snapshot”

`ioctl(phc_fd, PTP_SYS_OFFSET[_PRECISE])`

- Returns `(CLOCK_REALTIME, PTP_TIME)`
- Used to nudge **m** up or down over time

Piecewise-Linear Clock Model: $y=mx+c$



PCIe PTM IN ACTION

At t1', snapshot PTP Counter

Soon after t4', use t1, t4, t2', t3-t2 to compute:

- $linkDelay = \frac{[(t4-t1)-(t3-t2)]}{2}$
- **PTM Root time @ t1' = t2' - linkDelay**

Return the Cross-Timestamp:

- (PTM Root Time @t1', PTP Counter @t1')

Repeat as requested (by software)

Note: All PTM timestamps are in units of nanoseconds
 Note: Clock discontinuities are allowed at any time in the Upstream Port except between t1→t4, t1'→t4', etc.

SUMMARY

1. Time is distributed with precision around the globe today
2. Uncertainty in the last centimeters can be substantial and relevant
3. PCIe PTM is an optional, hardware feature that addresses the problem
4. Nanosecond-level accuracy from NIC to CPU, immune to system load
5. Silicon and Software support of PTM is available today



QUESTIONS?

BACKUP



SYSTEM CONFIGURATION

Measurements shown were performed using an Intel Atom® x6427FE Processor @ 1.90GHz, 4 cores, no hyperthreading, 4 MB cache, 8 GB of memory, PCIe Gen3

Ubuntu 18.04.6 LTS

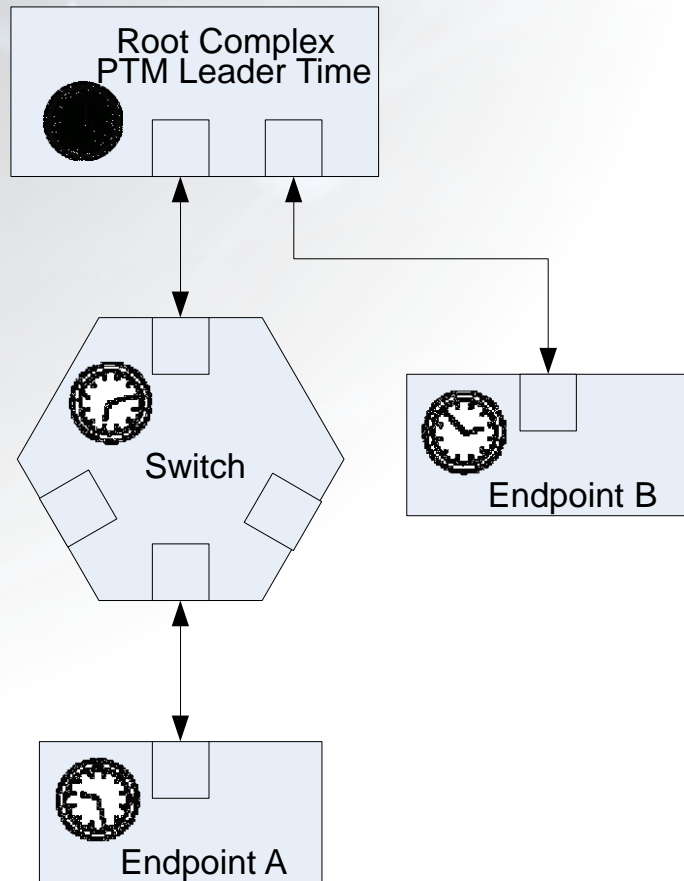
Linux Kernel version: 5.4.209 LTS

- Patched with <https://github.com/intel/linux-intel-lts>
 - Includes the TGPIO PPS driver

LinuxPTP version 3.1.1+

i225-LM Ethernet controller

PTM WITH A PCIE SWITCH



1. Endpoint A Initiates PTM Request **upward**
 - At the request of Software
 - Or on its own
2. Switch initiates PTM Request **upward**
3. Root Port Responds with the time **downward**
 - Delays measured and compensated
4. Switch Responds with the time **downward**
 - Delays measured and compensated
5. Endpoint knows the PTM Leader Time

Delays measured and added to the PTM Time

1. Link delays, PCIe SERDES Delays
2. Duration between 3. to 4. above