Microsoft

Recent Improvements to High-Accuracy Timekeeping in Microsoft Windows OS

vWSTS Conference - May 2020

Sarath Madakasira, Principal Software Engineer, Microsoft Corp.
Keith Mange, Principal Software Engineering Lead, Microsoft Corp.

Agenda

Timekeeping in Windows Server 2019

New PTP Features and Timekeeping improvements in Windows Server Next

PTP in Data Center Environment

Time sync test results in Server Next

Timestamping API, Windows Insider Program etc.

Timekeeping in Windows Server 2019

- Software Timestamping support in NTP Client/Server
- UTC-Compliant Leap Second support in Windows and NTP Client/Server.
  - https://doi.org/10.33012/2019.16762
- Framework for hardware timestamping of network packets
- Essential PTP Client(Slave) introduced
  - PTP over UDP/IPv4, Unicast/Hybrid modes using end-to-end Path Delays
  - Software Timestamping enabled and performs as well as NTP under similar conditions
  - https://github.com/microsoft/W32Time/tree/master/Precision%20Time%20Protocol
- Further reading:
  - https://techcommunity.microsoft.com/t5/networking-blog/evolution-of-timekeeping-in-windows/ba-p/778020

New PTP Features in Windows Server Next

- Hardware Timestamping integration with PTP Client
  - Currently restricted to 1 NIC on a machine
- Additional PTP Client features and modes, thanks in part to customer feedback.
  - Full-Multicast Mode (currently restricted to 1 NIC on a machine)
  - “Allow Any PTP Master” feature
    - Combined with multicast mode, enables plug-and-play PTP
  - PTP Leap Second Support
  - End-to-end residence time correction
  - Configurable PTP Domain Number

Timekeeping Improvements in Server Next

- Faster time convergence and improved clock stability when correcting larger time offsets
- Refining and adding to the dedicated W32time and PTP event logging for troubleshooting and record keeping.

PTP Data Center Deployment – a proof of concept

We deployed Windows PTP client in a data center environment as a proof of concept.

Learnings:

- Configuring existing network switches as boundary clocks allows scaling accurate time dissemination in a controlled environment – as was intended in the PTP specification.
- PTP + Software Timestamping in a v-cluster of about 40 machines in a large data center helped us realize time accuracies of ~50 microseconds across the cluster.

PTP Data Center Deployment – a proof of concept

Related info:

- Many models of network switches support some PTP modes, but not necessarily all modes.
  - Identifying the required modes upfront can help speed up the deployment.
- Network switches seem to have sufficiently accurate timekeeping to help us realize sub-50 microsecond accuracy across 3 switch hops.

Accuracy Testing Topology

Test topology consists two generic server blades connected via a generic switch (all a few years vintage).

We are running either an NTP or a PTP server on one of the machines and transferring the time information to an NTP or PTP Client on the other machine.

- The time source is the free running system clock on the server.
- Only new hardware components here are NICs that support Hardware Packet Timestamping supplied by Independent Hardware Vendors (IHVs).
- Relevant software includes Windows Server Next and NIC drivers supplied by IHVs.
- For testing NTP, we enabled Software Timestamping and used our Server Next NTP Server and Client (Not significantly different from the released versions).
- For testing PTP, we enabled Hardware Timestamping and used a private PTP Master along with our Server Next PTP Client.

NTP +SW Timestamping Test Results - Baseline

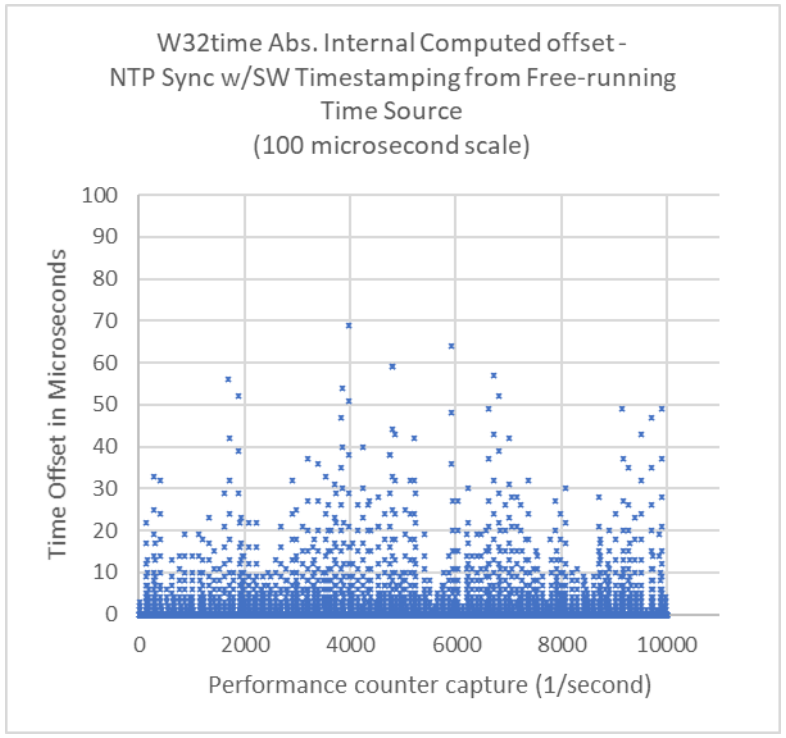

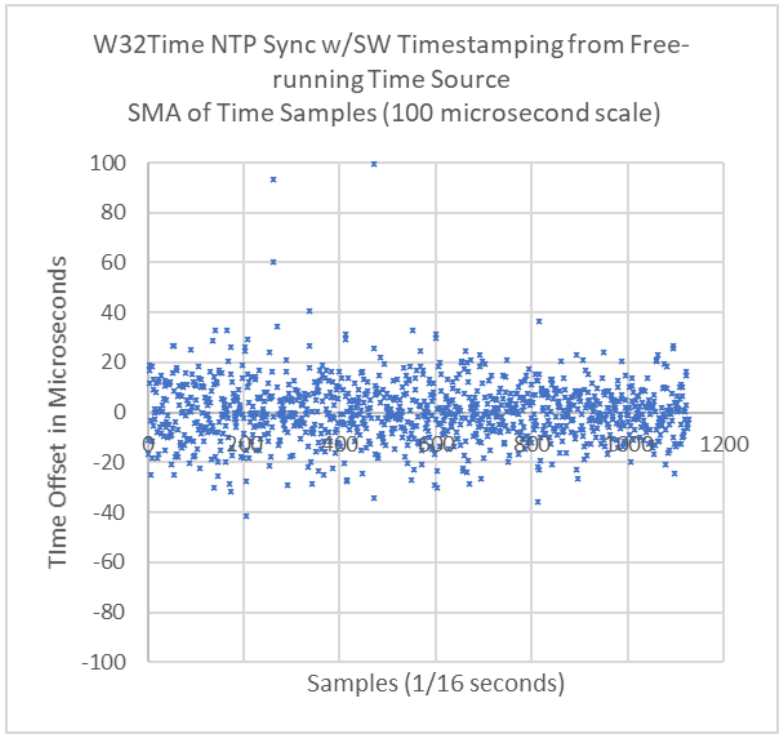

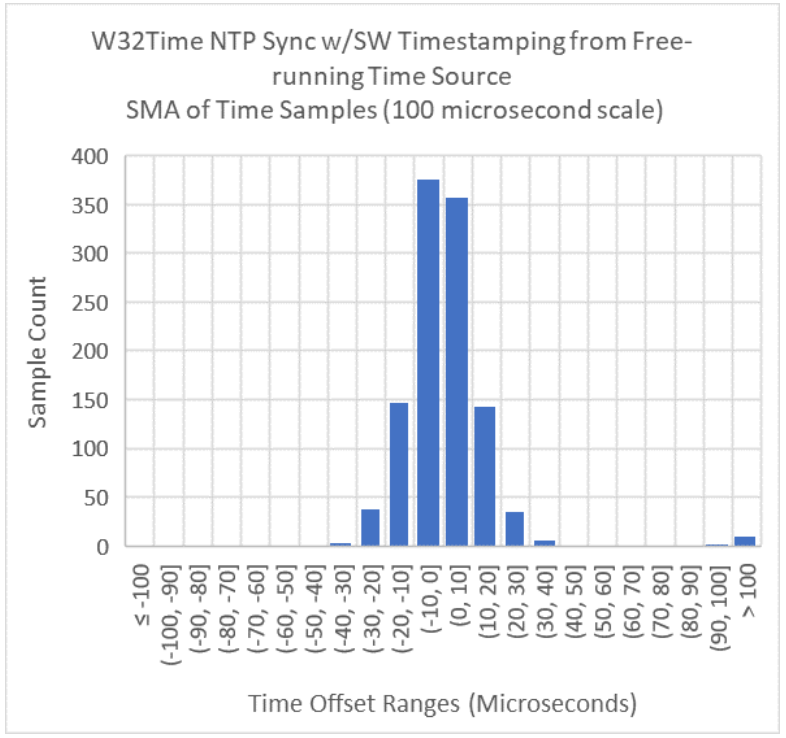

"Computed Time Offset" Performance Counter capture

NTP Sample Offset

NTP Sample Offset Distribution

PTP + HW Timestamping Test Results – Best Observed Scenario

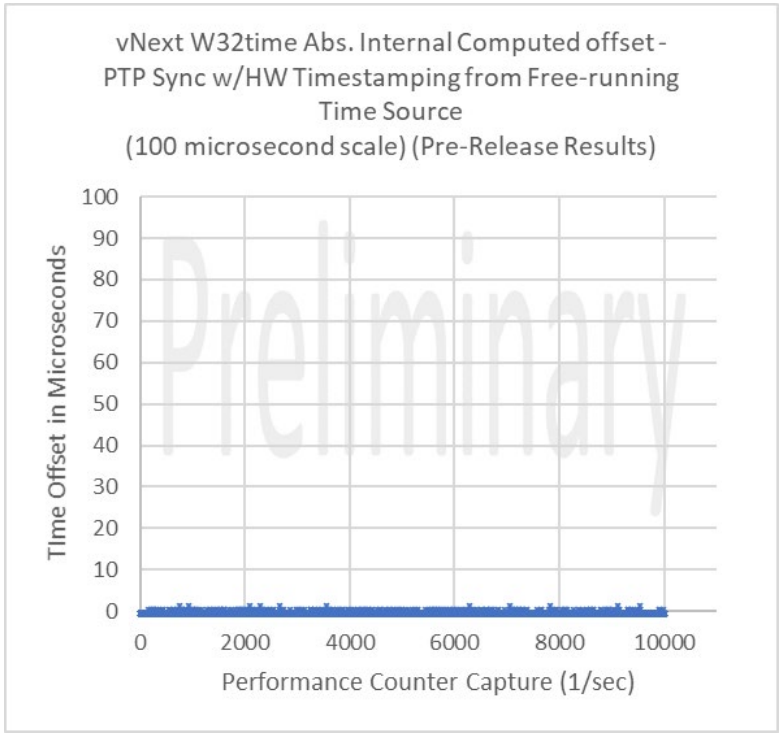

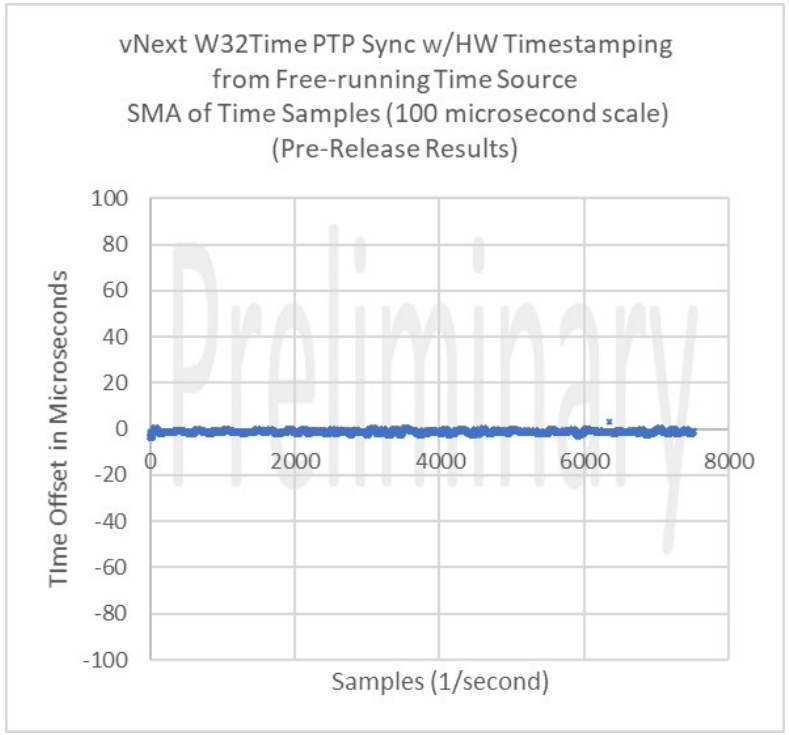

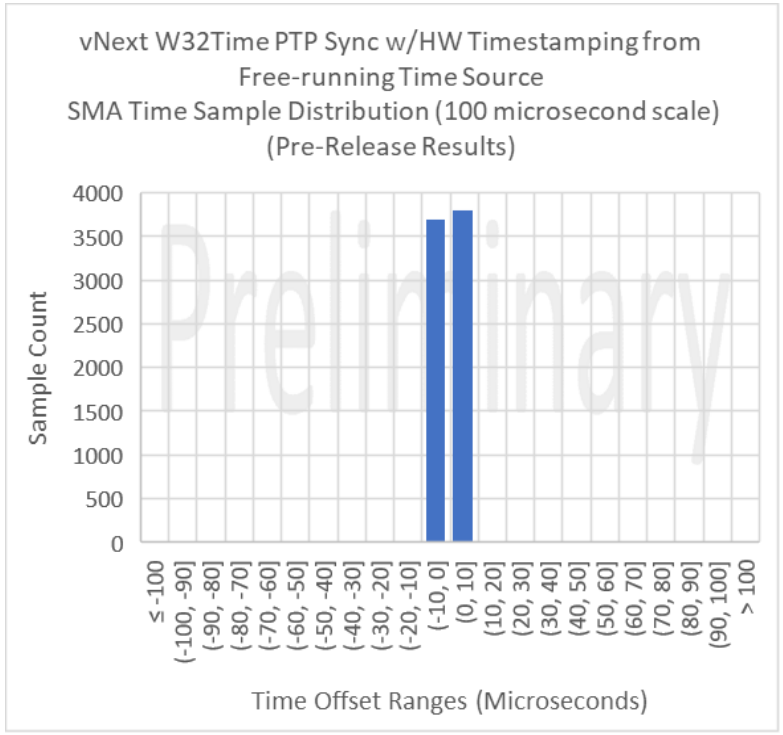

"Computed Time Offset" Performance Counter capture

PTP Sample Offset

PTP Sample Offset Distribution

NTP/PTP – Side-by-Side Comparison #1

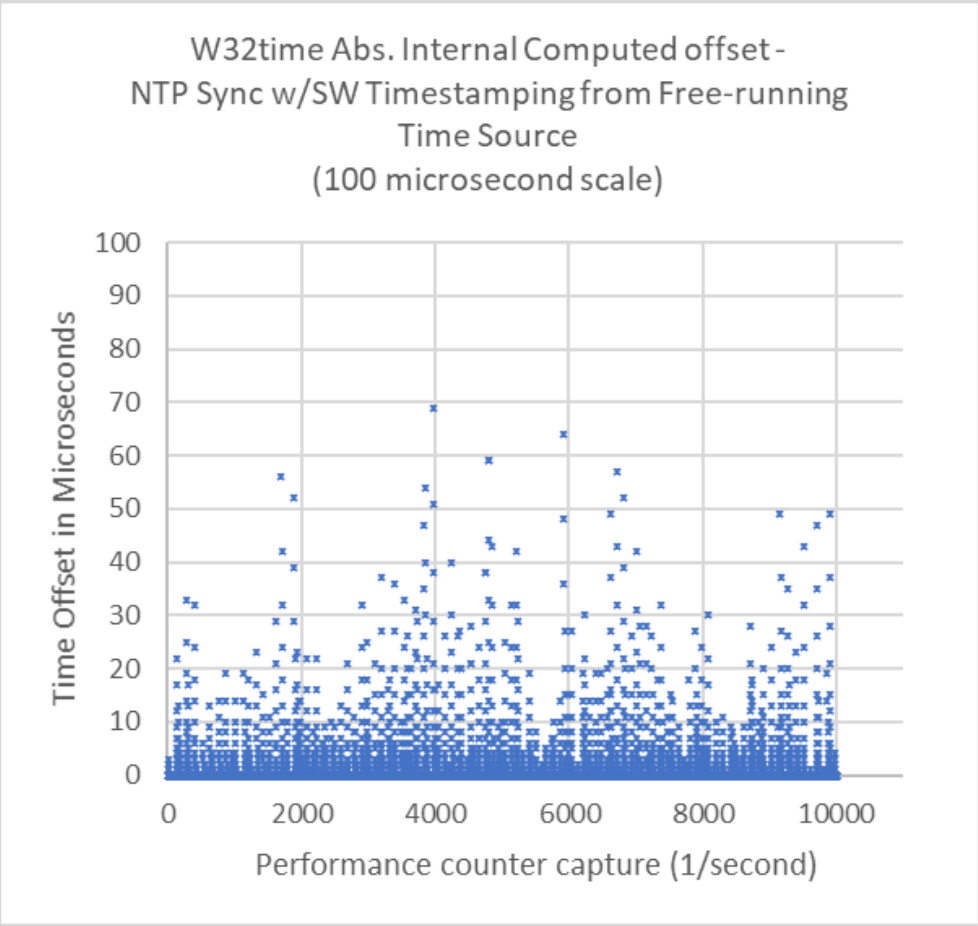

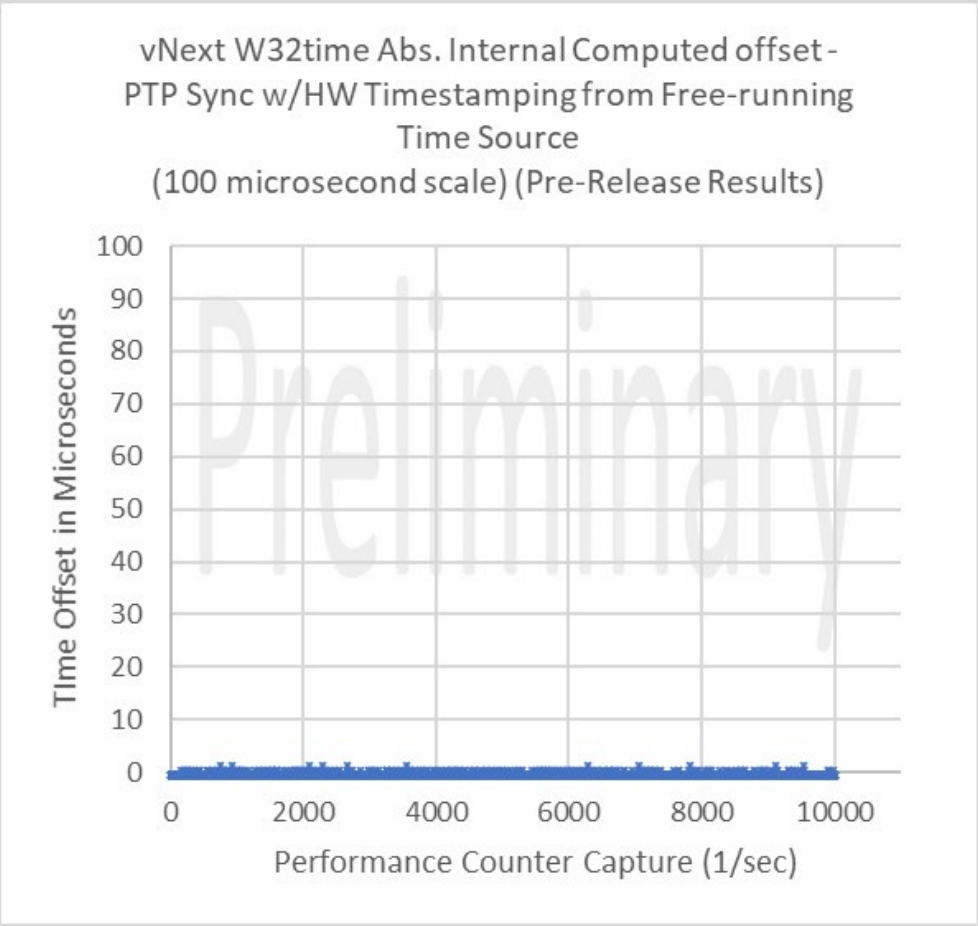

"Computed Time Offset" Performance Counter – Side by side

Improvements due to Hardware Timestamping integration with PTP

NTP/PTP – Side-by-Side Comparison #2

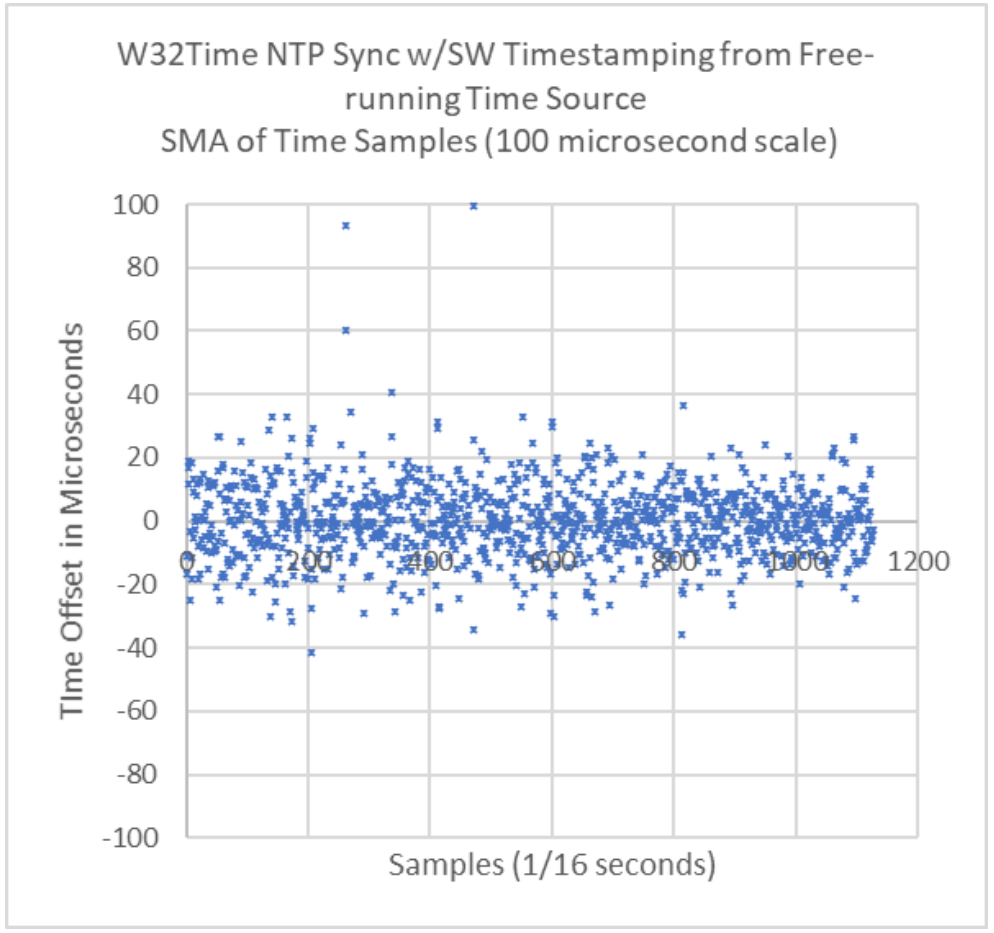

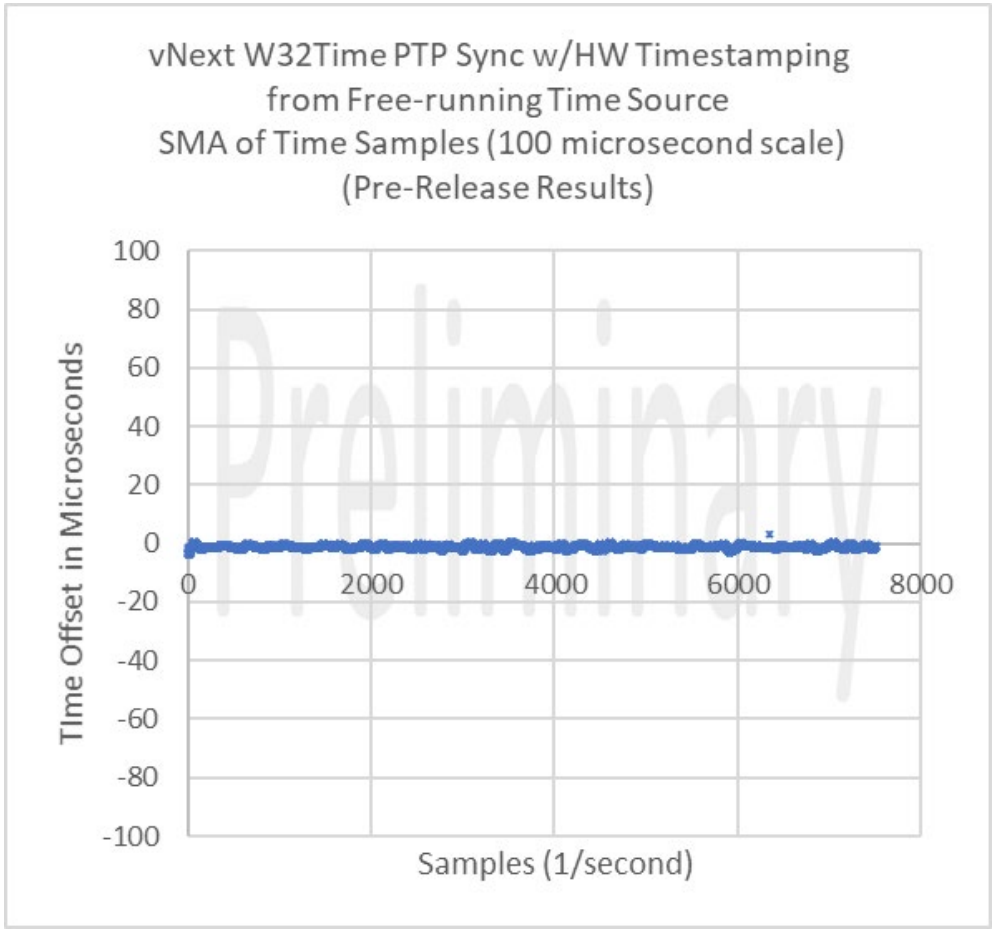

Time Sample Offset Comparison – Side-by-side

Improvements due to Hardware Timestamping integration with PTP

PTP + HW Timestamping Test Results – Closer look @ Best Observed Scenario

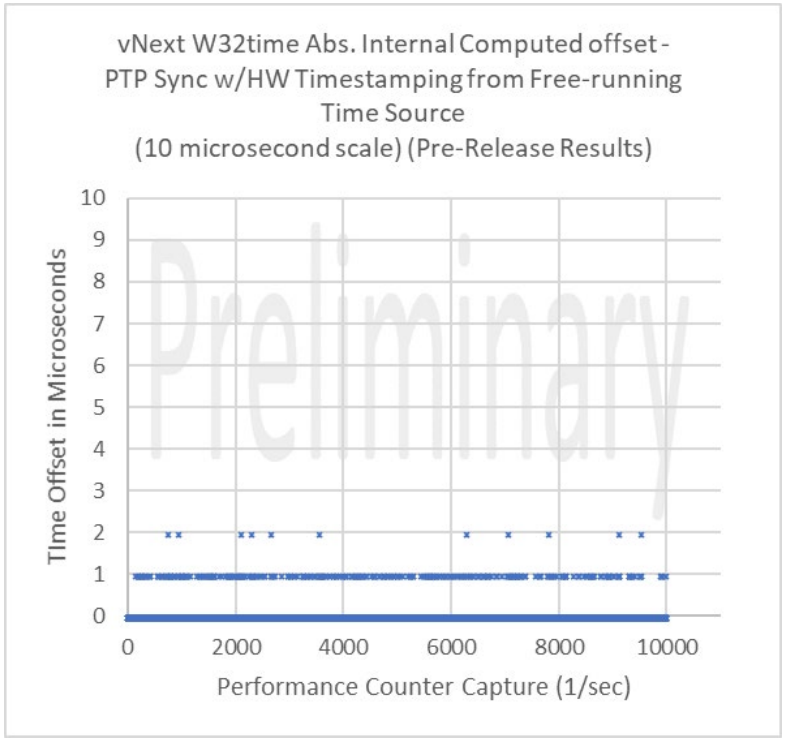

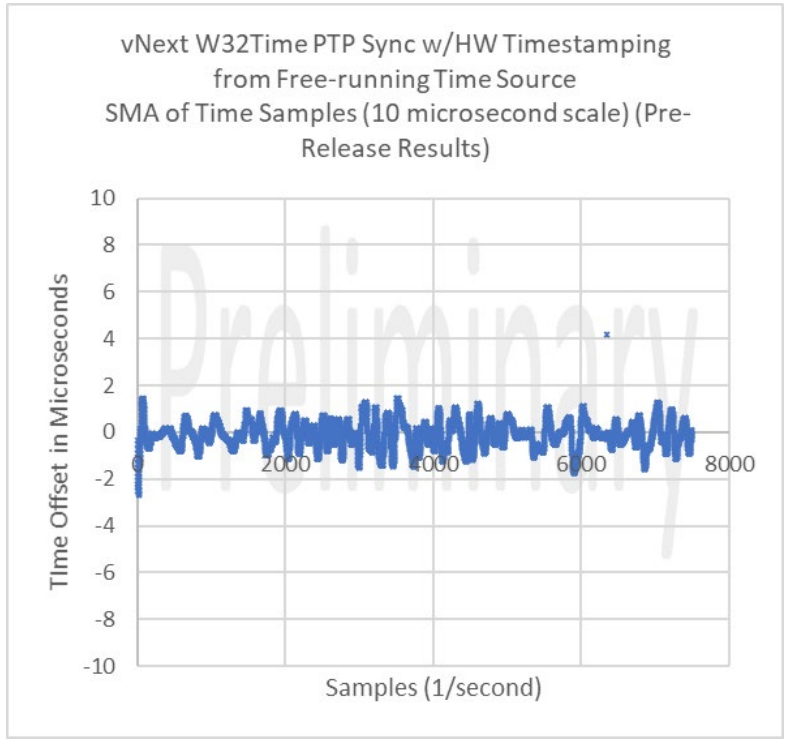

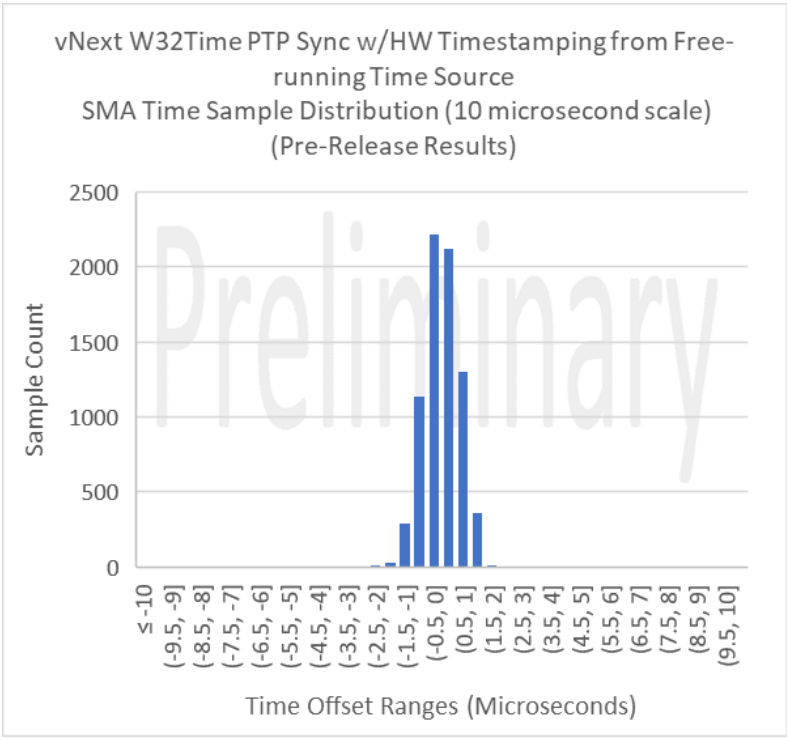

"Computed Time Offset"
Performance Counter capture

PTP Sample Offset

PTP Sample Offset Distribution

Showcases ±2 µs time offset – This is at least a Kessel Run away from Windows timekeeping from 10 years ago.

Timestamping for Windows Applications

- Windows API for accessing Timestamping functionality in Server Next should be released in due time.

Microsoft Windows Server Next

- Windows Server Next AKA pre-release Windows Server OS is available as part of the Windows Insider program.
  - Refer to the program for specifics: https://insider.windows.com/en-us/
- Time Synchronization is included in the Windows Client and Server OS.
- Stable features from pre-release OS are generally included in next major OS release.

Testing Windows Server Next

We normally would release results from topologies running released OS that could be replicated by our customers for their own time synchronization needs.

We are choosing to share the best results we have observed from Server Next because they have been compelling.

Work on Windows Server Next is still in progress and the time sync performance in the next Server release may be different than what is observed here.

Conclusion

We are excited about the possibility of awesome time sync accuracy in Server Next resulting from the integration of PTP with Hardware Timestamping.

We look forward to publishing detailed guidance and data for wider usage with the next Server release.

We are committed to making the best possible time synchronization widely available for everyone’s use.