



Maximizing Network Value

A Competitive Analysis of
Intel® Ethernet E830 Network Adapter
Efficiency and Timing Capabilities

AUTHOR

**Signal65 Network Performance
Analysis Team**

IN PARTNERSHIP WITH



FEBRUARY 2026

Executive Summary

In a global economy driven by high-speed data, automation, and regulations, the millisecond accuracy of legacy timing protocols like Network Time Protocol (NTP) is becoming increasingly outdated. Many businesses now depend on precise timing for critical applications across industries such as telecommunications, financial services, artificial intelligence (AI), manufacturing, and transportation. To address the shortcomings of legacy technologies, the Precision Time Protocol (PTP) standard provides sub-microsecond time synchronization. The selection of a PTP-capable Ethernet network adapter is not a commodity purchase but a foundational architectural decision with far-reaching implications for performance, compliance, and operational integrity across a variety of use cases.

This report presents a comprehensive performance evaluation of the Precision Time Protocol (PTP) for three leading Ethernet network adapters: the NVIDIA® MCX713104AS-ADAT ConnectX®-7 HHHL (NVIDIA CX7), the Broadcom® BCM957504-P425G (Broadcom P425G), and the Intel® Ethernet Network Adapter E830-XXVDA2 (Intel E830-XXVDA2). References to the Intel E830 are for the family of network adapter products. Specific testing results, however, were derived from the E830-XXVDA2 SKU.

The analysis quantifies and compares network adapters based on time accuracy, stability, robustness, power efficiency, flexible and broad port density options, and other critical selection factors, to identify the optimal hardware solution for mission-critical applications.

Key Findings

The Intel E830-XXVDA2 demonstrated superior PTP accuracy across all critical metrics and showed significantly better resilience to network congestion. While all three network adapters provided sub-microsecond timing accuracy:

- The Intel E830-XXVDA2 demonstrated the highest stability under network congestion, with jitter increasing by only 4.1% when subjected to load;
- The NVIDIA CX7 delivered strong baseline performance but exhibited up to 43.4% more timing error rates under load;
- And the Broadcom P425G exhibited the least resilience to network traffic, with jitter increasing by over 150% under the loaded test scenario.

Based on the test results, the Intel E830 is recommended for timing-critical deployments. Its ability to maintain exceptional stability and accuracy irrespective of network load results, coupled with its engineering for Xeon compatibility, indicates a superior PTP and hardware architecture. This makes it the lowest-risk and most future-proof foundation for enterprise timing, capable of meeting the demands of today's applications and the more stringent regulatory requirements.

The Enterprise Imperative for Precision Timing

For decades, the Network Time Protocol (NTP) has been widely utilized, providing millisecond-level accuracy sufficient for the ordering of system log files and basic transaction records. This level of precision was adequate in an era where operational timescales were driven by human perception and loosely synchronized systems were sufficient. However, modern data processing, real-time applications, and the proliferation of distributed, data-intensive systems have reduced these timescales to the nanosecond scale. This has created a critical gap between the capabilities of legacy timing protocols and the requirements of modern enterprise networks.

- **Telecommunications:** Inaccurate timing can lead to signal interference, reduced quality of service, and the failure of low-latency services, potentially resulting in Service Level Agreement (SLA) breaches.
- **Financial Markets:** Microsecond ambiguity can compromise the accurate sequencing of transactions, posing a risk to market transparency and compliance with regulatory mandates (e.g., MiFID II).
- **Industrial Automation:** Timing misalignment can cause unsynchronized actions in robotics and machinery, leading to production errors, equipment damage, and operational downtime.
- **AI/ML Clusters:** Precise synchronization is essential for aligning sensor data and maintaining the deterministic order of inter-node communication required by high-performance distributed machine learning models.

The reliance on older timing technologies in these environments has shifted from a standard practice to a significant business liability—jeopardizing regulatory compliance, operational efficiency, and bottom-line profitability—and ultimately undermining the viability of next-generation technologies.

Comparative Performance Analysis

Signal65 conducted a series of tests on the latest Intel enterprise network adapter, the Intel E830-XXVDA2, along with two leading competing network adapters. Throughout the report, we refer to these two competitors' cards as Broadcom P425G and NVIDIA CX7.

Note: *This analysis utilized standard, general-purpose network adapters (the Intel E830-XXVDA2, Broadcom BCM957504-P425G, and NVIDIA ConnectX-7 MCX713104AS-ADAT) to ensure broad applicability across all deployment environments, from cloud to edge.*

Testing was performed in a lab environment that contained a Calnex Paragon-neo NRZ system as the primary instrument for generating and measuring PTP timing accuracy. The network adapters were tested inside of an enterprise server, using the Intel® Xeon® 6 processor. Details of this equipment are provided in Appendix C.

Key Performance Metrics for PTP Analysis

To provide a rigorous and meaningful comparison of the network adapters, this report focuses on three fundamental metrics that define the quality of a PTP clock:

- **Time Error (Offset):** This is the instantaneous difference between the Receiver clock time and the Transmitter clock time. While a stable error rate can be compensated for using software, achieving lower time error values is preferable for maintaining high precision.

- **Jitter (Standard Deviation):** Jitter quantifies the variability in time error over a specified measurement period. It is expressed as the standard deviation of the time error samples, indicating the consistency of the clock's performance. Lower jitter values suggest more stable and reliable time synchronization.
- **Peak-to-Peak Error (Max-Min):** This metric measures the absolute range between the maximum and minimum time errors recorded during the test. It provides insight into the extent of time error fluctuations, with smaller peak-to-peak values indicating tighter control over timing variations.

Test Methodology

The results of this performance evaluation are based on data from a controlled test environment equipped with one of the target network adapters (Broadcom P425G, NVIDIA CX7 and Intel E830-XXVDA2) configured as a PTP time Receiver. These time Receivers were synchronized to a common, high-precision PTP grandmaster clock originating within a specialized test device, a Calnex Paragon-neo NRZ, which serves as the authoritative time reference for the test network. Each network adapter was subjected to two distinct, eight-hour test scenarios:

- **No-Load Condition:** This test established a baseline performance profile with minimal background network traffic, isolating the inherent capabilities of the network adapter's hardware and PTP software stack.
- **Loaded Condition:** This test introduced significant network traffic to simulate a congested, real-world environment. This scenario is designed to stress the PTP implementation and evaluate its robustness against packet delay variation.

The key performance indicators are outlined in the next section, with detailed information provided in the appendices.

Performance Results

Signal65 analyzed both baseline performance of each network adapter under ideal, no-load conditions along with testing during network traffic load. The baseline, no-load helps isolate the inherent capabilities of the hardware and driver implementations from the variable effects of network congestion, establishing a benchmark for stability and accuracy. Network congestion is a general term meant to indicate significant network traffic conditions exist.

Mean Time Error

The mean time error results from our experiments for each network adapter are shown below in Figure 1.

No Load Test

The mean time error (MTE) quantifies the average deviation of a system clock from the reference time provided by the network grandmaster clock over a specified length of time. The MTE is calculated by averaging the time offsets recorded at regular intervals between the system clock and the grandmaster clock. Under ideal conditions, a well-calibrated system should converge to a mean error close to zero. The measured mean errors were:

Network Adapter	Intel E830-XXVDA2	Broadcom P425G	NVIDIA CX7
Jitter MTE (ns)	-7.3	-72.1	76.7

The Intel E830-XXVDA2 demonstrates a superior out-of-the-box performance, settling at a mean offset of only -7.3 ns. In contrast, both the NVIDIA CX7 and Broadcom P425G network adapters stabilized with significant systemic error offsets of +76.7 ns and -72.1 ns, respectively. While a known constant offset may be compensated for, this would have to be benchmarked per system configuration.

Load Test

Under network load, packet delays can introduce a timing bias if not properly handled. The change in mean time error from the no-load condition is a key indicator of this effect. The measured mean errors under load were:

Network Adapter	Intel E830-XXVDA2	Broadcom P425G	NVIDIA CX7
Jitter MTE (ns)	-7.3	-99.4	101.5

The Intel E830-XXVDA2's performance is exceptional; its mean error remained virtually unchanged from the no-load condition, demonstrating an almost complete immunity to load-induced offset bias. Both the NVIDIA CX7 and Broadcom P425G adapters, however, showed a shift in their mean error. NVIDIA CX7's offset increased from 76.7 ns to 101.5 ns, while Broadcom P425G's offset increased from -72.1 ns to -99.4 ns. This indicates that these two cards were less effective at mitigating the systematic delay introduced by network congestion, allowing it to negatively impact their overall accuracy.

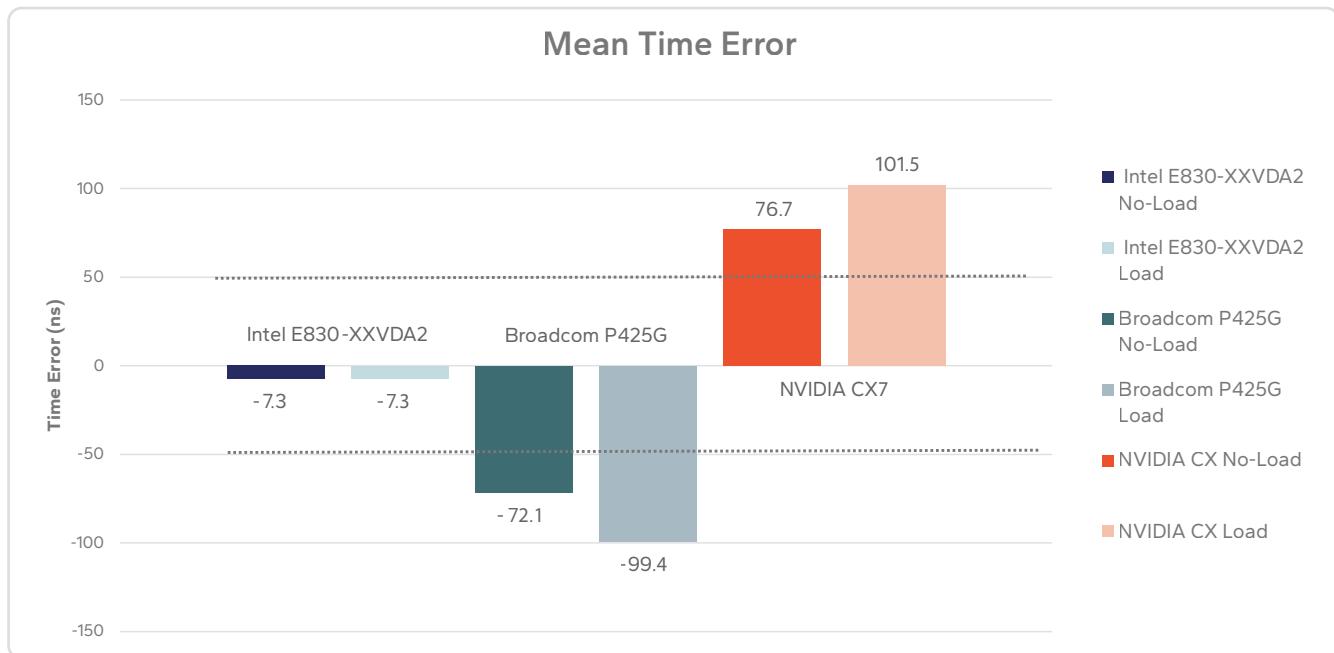


Figure 1: Average Mean Time Error (Source Signal65)

Jitter and Stability

As mentioned earlier, jitter measures the variability of the clock in question, with lower jitter values indicating a more stable and reliable time synchronization. The standard deviation for each network adapter under no-load conditions was:

Network Adapter	Intel E830-XXVDA2	Broadcom P425G	NVIDIA CX7
Jitter (ns)	12.4	17.9	15.2

The data indicates that the Intel E830-XXVDA2 is the most stable clock at baseline, exhibiting the lowest jitter at 12.4ns followed by the NVIDIA CX7 and Broadcom P425G with higher standard deviations. These results establish a clear hierarchy of baseline stability, with the Intel E830-XXVDA2 leading the comparison.

The standard deviation for each network adapter under loaded conditions was:

Network Adapter	Intel E830-XXVDA2	Broadcom P425G	NVIDIA CX7
Jitter (ns)	12.6	45.6	21.8

The performance divergence among the cards becomes apparent under load. The Intel E830-XXVDA2 continues to exhibit outstanding stability, with its jitter increasing by a mere 0.5ns. The NVIDIA CX7 shows a respectable level of resilience, with its jitter increasing by approximately 6.6ns. The Broadcom P425G, however, demonstrates a significant degradation in performance with its jitter more than doubling, reaching 45.6ns.

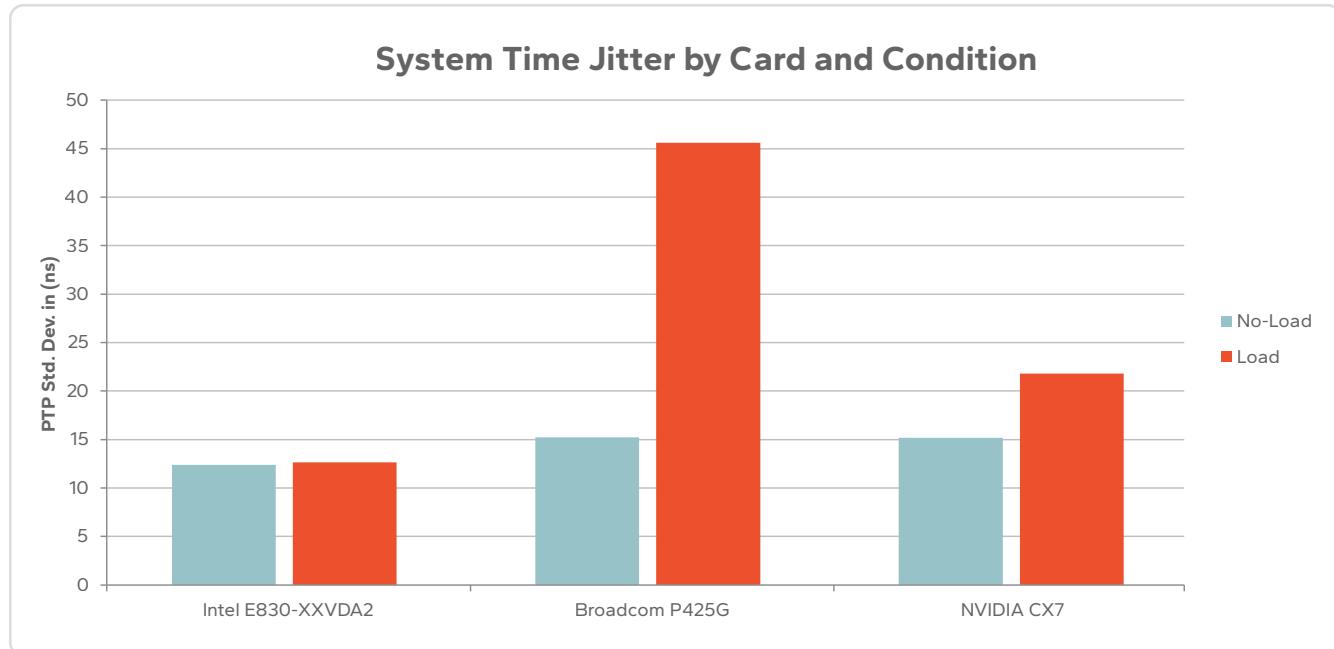


Figure 2: Standard Deviation of Jitter Rates (Source Signal65)

Peak-to-Peak Error and Outlier Analysis

Peak-to-Peak (P-P) error measures the absolute range between the minimum and maximum observed time errors, defining the worst-case performance envelope. This metric is particularly important for identifying the presence of significant, infrequent outliers that may not be represented by the average and standard deviation alone. This can lead to SLA violations even when the average values appear within specifications.

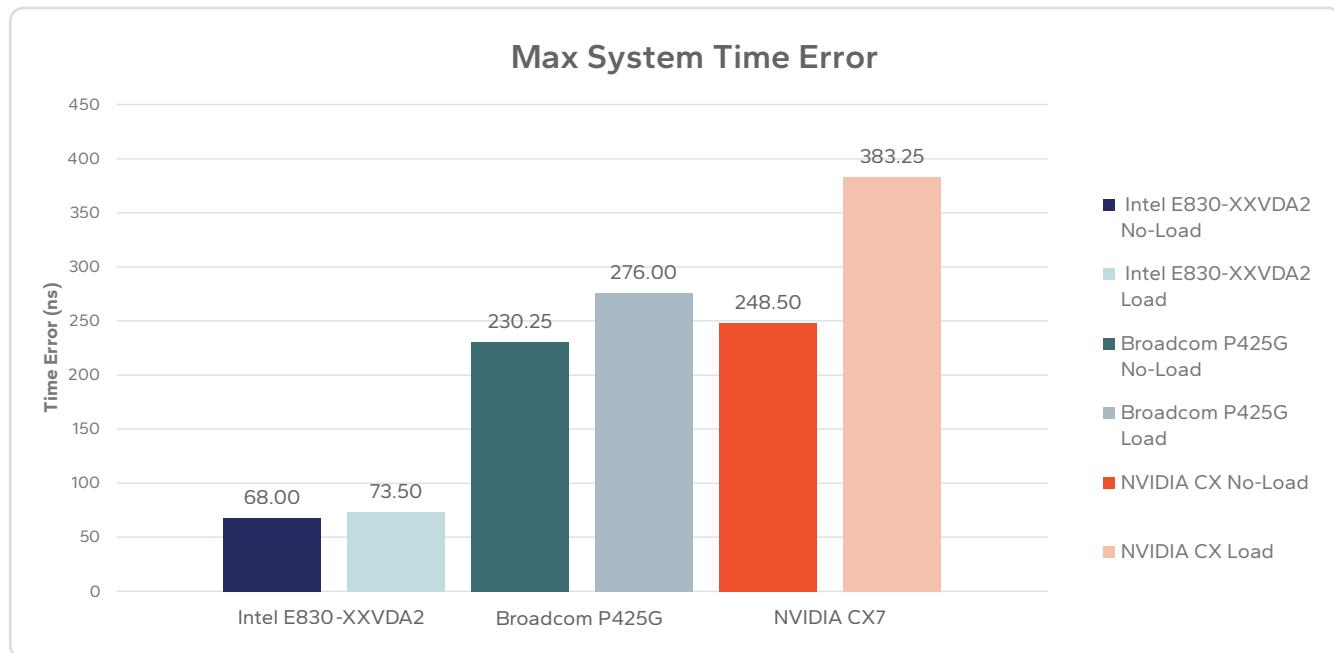


Figure 3: System Time Error Rates (Source Signal65)

The relationship between jitter and P-P error provides a deeper understanding of each network adapter's behavior. The Intel E830-XXVDA2's tight P-P range of 68.0ns is consistent with its low standard deviation, indicating that its time errors are not only small on average but also well-bounded. Conversely, the NVIDIA CX7's large P-P error of 248.5ns is disproportionately high relative to its low standard deviation of 15.2ns. This combination indicates that while most of the NVIDIA CX7's time error measurements are tightly clustered around its mean, the network adapter is susceptible to infrequent but significant outlier events, even in the absence of network load. For applications that are intolerant of any single erroneous timestamp, such as high-frequency trading or precise industrial control, this behavior could be a critical concern. The Broadcom P425G card exhibits a moderate P-P error that is more consistent with its standard deviation.

Quantifying Performance Degradation Under Load

To directly compare the robustness of each network adapter, the percentage increase in standard deviation from the no-load to the loaded condition was calculated. This metric normalizes the results and provides the clearest indication of how well each implementation handles real-world network conditions.

Metric	Intel E830-XXVDA2	Broadcom P425G	NVIDIA CX7
Std. Dev. (No-Load, ns)	12.1	17.9	15.2
Std. Dev. (Loaded, ns)	12.6	45.6	21.8
Degradation (Jitter Increase)	4.1%	154.7%	43.4%

Table 1: Jitter Degradation Under Network Load

Platform Time Synchronization

The definitive test of a network adapter's PTP implementation is its ability to maintain timing stability under network load, which separates robust, predictable hardware from implementations that degrade under real-world conditions. Under ideal conditions, all three network adapters delivered good baseline stability. However, the Intel E830-XXVDA2 showed superior out-of-the-box performance with a near-zero mean time error. Critically, the NVIDIA CX7 network adapter revealed that despite its low average jitter, its disproportionately large P2P error range indicated a susceptibility to infrequent but significant outlier events, a key concern for applications with strict timing boundaries.

The performance divergence became stark under network congestion. The Intel E830-XXVDA2's performance remained virtually unchanged, demonstrating exceptional immunity to network-induced jitter. In contrast, the NVIDIA CX7 performance degraded moderately, and the Broadcom P425G's jitter more than doubled, indicating it is unable to cope with realistic network conditions. This distinction is critical, as the true value of a timing solution is its predictability in a dynamic production environment. The Intel E830-XXVDA2 provides deterministic timing integrity, while the others introduce a level of uncertainty that reframes the hardware selection process from a simple spec-sheet comparison to a crucial exercise in risk management.



Value of Integrated Platform Architecture

Intel Platform Integration for Optimized Performance

The test results demonstrating the superior stability of the Intel E830-XXVDA2 are linked to **Intel's platform-level integration strategy**. When an Intel network adapter is paired with Intel® Xeon® 6 processor, it creates a deeply integrated system where components are engineered to work in concert. This synergy unlocks a level of optimization in timing and data processing that third-party components operating in the same server cannot replicate.

Integrating Precision Timing: A Unified System Approach

Achieving precision timing involves more than just obtaining accurate time from the network; it requires delivering this time to the operating system and end-user application without degradation. Intel's architecture addresses this through a unified timing solution built on three key technologies:

- **PTP Hardware Support:** The Intel E830-XXVDA2 features hardware support for the Precision Time Protocol (PTP), allowing it to timestamp incoming and outgoing packets with nanosecond precision. This capability ensures the precise capture of PTP events timings, minimizing the variability introduced by software processing delays.
- **Precision Time Measurement (PTM):** PTM extends the network adapter's nanosecond-level network accuracy to the host system's clock which can in turn be utilized by applications. As a component of the PCI Express specification, PTM employs a hardware-based "cross-timestamping" mechanism to precisely calculate and correct for latency and jitter as the time signal traverses the PCIe bus. Without this hardware-based correlation, applications must rely on slower, software-based methods to access the network adapter's clock, a process that can introduce microseconds of timing uncertainty, reducing the accuracy of timing protocols. PTM ensures the application receives a network-synchronized time value with sub-microsecond accuracy. Competing network adapters lacking PTM support are forced to use these less precise software methods, creating an architectural disadvantage in timing accuracy. Without PTM, software access to the network adapter's time would introduce significant jitter, ranging from hundreds of nanoseconds to microseconds, thereby undermining the benefits of PTP.
- **Time-Aware General Purpose I/O (TGPIO):** These general-purpose I/O pins are linked to the host CPU's hardware clock, enabling both the import/export of the unified time domain to/beyond the server and the validation of time at the host as well as. They can be used to import a Pulse Per Second (PPS) signal to discipline the system clock and generate a precise PPS output to synchronize other devices such as cameras and sensors that lack PTP capabilities. Additionally, TGPIO serves as a key proof point for verifying system-level timing accuracy when coupled with a signal analyzer.

Accelerated Data Movement and Processing

Accurate time is only useful if the system can act on time-sensitive data with minimal latency. Intel's platform-level technologies ensure that data processing is as efficient as time synchronization through multiple instances of software and hardware co-optimization:

- **Intel® Data Direct I/O (DDIO):** A cornerstone of Intel's platform advantage, DDIO allows the E830 network adapters to write incoming network data directly into the Xeon processor's last-level cache. This bypasses the traditional, higher-latency path to system memory (RAM), making data immediately available to CPU cores. For latency-sensitive applications like high-frequency trading, this dramatically reduces processing delays. Data Plane Development Kit/Vector Packet Processing (DPDK/VPP) software optimizations result in more efficient cache utilization and improved system performance, especially when paired with Intel Xeon processors, ensuring superior efficiency and performance.
- **Dynamic Device Personalization (DDP):** DDP enables dynamic reconfiguration of the Intel Ethernet E830 network adapter's packet-processing pipeline to match specific workload requirements, allowing systems to optimize traffic classification and handling strategies for maximum efficiency across diverse network applications (e.g., leveraging AVX instructions for processing multiple packets per cycle). This allows for improved efficiency by offloading packet classification and increased throughput in minimizing bottlenecks.
- **Intel® QuickAssist Technology (QAT):** The architecture is further enhanced by deep software and hardware co-optimization. For example, deep integration with the DPDK and the ability to offload tasks like cryptography to the Intel Xeon's built-in QAT free up CPU cores for primary application logic, boosting overall throughput.

When Intel Xeon and Intel Ethernet E830 Network Adapter operate together, timing, data, and compute converge into a single synchronized platform. The result is deterministic precision, validated in hardware, that competitors can't replicate through software or standalone network adapter design.

Economic and Power Consumption Considerations

Acquisition Cost Comparison

Beyond technical considerations, economic factors, particularly acquisition costs, play an important role in selecting foundational components. There are many different configuration options available for network adapters with differing port densities and speed. For the cards analyzed (NVIDIA, Broadcom, and Intel), standard port density options typically include dual and quad ports. Intel uniquely extends this offering with a physical 8-port option in its newly launched E830-XXVDA8F OCP 3.0 DSFF adapter. This high-density design is specifically tailored for saving hardware CAPEX by reducing the overall number of components needed in consolidated cloud and edge architectures.

To make the comparison fair, we have chosen to utilize only 25Gb connection speeds, and evaluate cost on a per port basis, making the port density options less significant. As shown below, the Intel E830 options are significantly less expensive than the competing options on a per port basis.

Network Adapter Model	Port Speed	Number of Ports	Total Cost (\$)	Cost per Port (\$/port)	% Higher than Intel E830 (2-port)
Intel E830-XXVDA2	25 Gbps	2	\$216 (Suggested)	\$108.00	0%
Intel E830-XXVDA8F (OCP 3.0)	25 Gbps	8	\$976 (Suggested)	\$122.00	13.0%
NVIDIA ConnectX-6	25 Gbps	2	\$389 (Source)	\$194.50	80.0%
NVIDIA ConnectX-7	25 Gbps	4	\$962 (Source)	\$240.50	122.7%
Broadcom BCM957414A4142CC	25 Gbps	2	\$269 (Source)	\$134.50	24.5%
Broadcom BCM957504-P425G	25 Gbps	4	\$635 (Source)	\$158.75	46.3%

Note: The Intel E830-XXVDA8F (8-port OCP model) carries a slightly higher price per port than the Intel dual-port card. This is offset by its great I/O density, which achieves the capacity of four dual-port cards in a single slot and conserves precious PCIe resources for essential components like GPUs or high-end storage.

Power Utilization Analysis

Another consideration is power utilization and its implied environmental impacts and associated operational costs. The following table outlines the power consumption of relevant competitive products currently available on the market:

Network Adapter Model	Total Ports	Total Card Power (Typical/Operational)	Est. Power Per Port (W)	Notes
Intel E830-XXVDA2	2	8.8W (Typical, DAC)	4.4 W	Power specified as Typical
Intel E830-XXVDA8F (OCP 3.0)	8	15.5W (Typical, DAC)	1.94 W	Power specified as Typical
NVIDIA ConnectX-6	2	11.53 W (Typical, DAC)	5.77 W	Power specified as Typical
NVIDIA ConnectX-7	4	15.8W (Operational)	3.95 W	Power specified as Operational
Broadcom BCM957414A4142CC	2	10.9W (Typical, DAC)	5.45W	Power specified as Typical
Broadcom BCM957504-P425G	4	16.9W (Operational)	4.23 W	Power specified as Operational

Based on the published power efficiency data available for the competitive products, the Intel E830 product line offers competitive power efficiency on top of the lower acquisition cost and superior performance demonstrated by the test results presented above.

Note: The 8-port Intel Ethernet E830-XXVDA8F for OCP 3.0 DSFF commands a higher price per port than the Intel dual-port card because its core benefit is I/O density optimization. It achieves the capacity of four dual-port cards in a single slot, conserving PCIe resources for essential components like GPUs or high-speed storage. Notably, Intel's high-density solution often achieves a lower price per port relative to comparable high-port-count solutions from its competition, providing better value in specialized, resource-constrained server deployments.

Analysis and Strategic Conclusions

Based on the combination of superior stability under network congestion, the architectural advantage of Precision Time Measurement (PTM) for host time delivery, and competitive total cost of ownership (TCO) considerations, the Intel E830 network adapter is the demonstrated optimal solution for timing-critical enterprise deployments.

Meeting business requirements in the following environments requires high PTP fidelity:

- Ultra-High Precision Requirements (e.g., 5G Telecommunications): For applications demanding Class B timing accuracy (maximum Time Error of nanoseconds), the Intel E830-XXVDA2 was the only tested network adapter that met the stringent requirement, recording a worst-case error under load of approximately 40 ns. The other adapters exceeded this tolerance level.
- Microsecond Tolerance Requirements (e.g., Industrial Automation, Finance): For applications with sub-microsecond tolerances, both the Intel and NVIDIA CX7 network adapters performed within the acceptable range under baseline conditions. However, the demonstrated instability of the Broadcom P425G under load saw jitter increasing by over 150%, indicating a significant risk of unpredictable timing in congested network environments. This variability compromises its suitability for deployments requiring verifiable, consistent compliance, such as financial transaction logging.

The test results confirm that PTP hardware selection requires validation under realistic, congested network load conditions. Relying solely on baseline performance or datasheet specifications does not adequately assess the risk of timing degradation in operational environments. Investment in a robust PTP infrastructure, exemplified by the deterministic stability and integrated architecture of the Intel E830, minimizes risk and provides a performance-validated foundation for modern synchronized systems.



Appendix A

Products Tested:

- Intel E830-XXVDA2
- Broadcom BCM957504-P425G
- NVIDIA ConnectX-7 MCX713104AS-ADAT

Note: It is important to understand that NVIDIA has a Telecom – 5G specific CX7 card with enhanced PTP timing. We did not utilize this specialized card during testing, instead using standard, general-purpose network adapters to understand the impact of timing outside of the telecom industry.

Detailed PTP Metrics and Industry Tolerance Levels

Key Timing Metrics

To properly quantify PTP performance, it is essential to understand the following, previously outlined in the paper:

- **Time Error (Offset):** This is the instantaneous difference between the Receiver clock time and the Transmitter clock time.
- **Jitter (Standard Deviation):** Jitter quantifies the variability in time error over a specified measurement period.
- **Peak-to-Peak Error (Max-Min):** This metric measures the absolute range between the maximum and minimum time errors recorded during the test.

ITU-T G.8273.2 Telecom Boundary Clock (T-BC) Classes

The telecommunications industry has the most detailed specifications for PTP performance, defining different classes of "Telecom Boundary Clocks" (T-BCs) with nanosecond-level precision to support demanding 5G applications. The noise generation limits for these clocks are as follows:

Class	Maximum Time Error (Max TE)	Constant Time Error (cTE)	Primary Application
Class A	100 ns	50 ns	General 5G backhaul
Class B	70 ns	20 ns	Enhanced 5G backhaul
Class C	30 ns	10 ns	5G fronthaul applications
Class D	5 ns (low-pass filtered)	N/A	Most demanding fronthaul networks

Appendix B: Time Series Graphs

Data Sources

The analysis utilizes two sets of files for each of the six test runs:

- **Raw Data (CSV Files):** These files contain the high-resolution time series data, with one time error measurement captured per second for the duration of each 8-hour test. This granular data, comprising approximately 28,800 data points per test, is essential for analyzing long-term drift and observing the detailed character of the time error.
- **Summary Statistics (PDF Reports):** The accompanying PDF reports provide pre-calculated summary statistics for most of the test runs, including key metrics such as Mean Time Error, Minimum Error, Maximum Error, and Standard Deviation. These values provide a high-level overview of the performance in each scenario.

Data Processing Note

To maintain a complete and consistent comparative dataset, the primary statistical metrics (Mean, Standard Deviation, Min, Max) for these scenarios were calculated directly from the raw data contained in their respective CSV file. All other data points are taken from their corresponding summary reports.

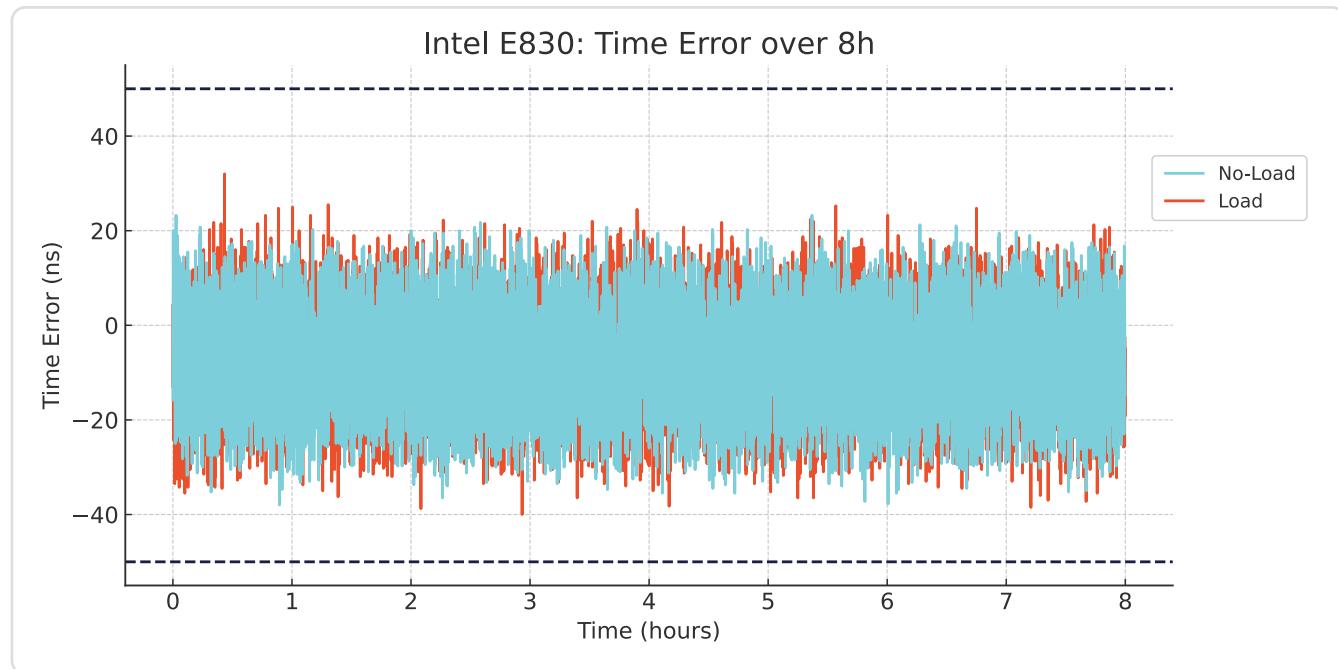


Figure B-1: Intel E830-XXVDA2 - Error Rates Over 8hr Test (Source Signal65)

Broadcom P425G: Time Error over 8h

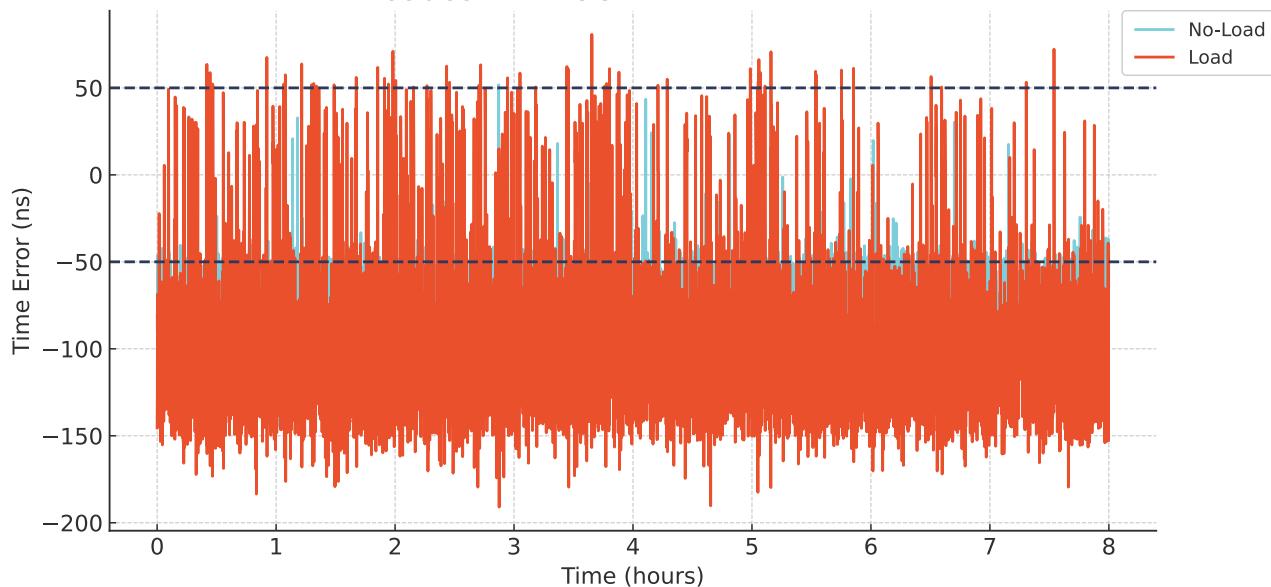


Figure B-2: Broadcom P425G - Error Rates Over 8hr Test (Source Signal65)

NVIDIA CX-7: Time Error over 8h

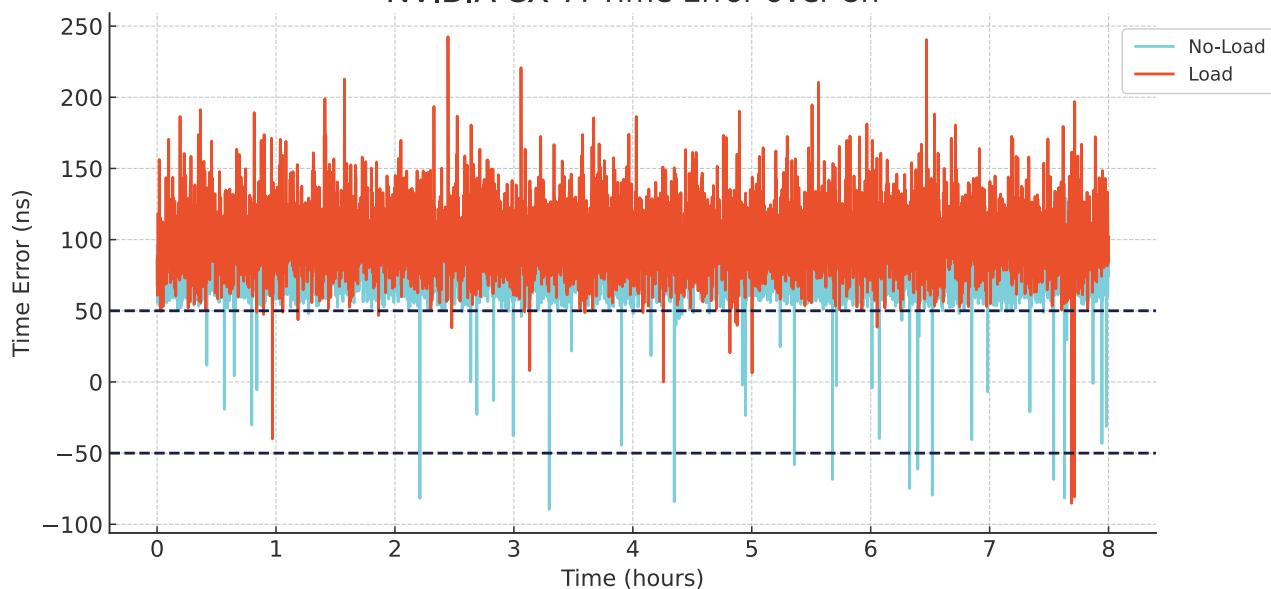


Figure B-3: NVIDIA CX7 - Error Rates Over 8hr Test (Source Signal65)

Appendix C: System & Software Information

System Configuration

Hardware		
Motherboard		Kaseyville RP
CPU	Product	Intel® Xeon® 6556P-B
	Speed (MHz)	2.3 GHz (base) - 3.5 GHz (max)
	No of CPU	36
	LLC Cache	144M
Chipset		Kaseyville
Memory	Vendor	Hynix
	Type	DDR5 RDIMM
	Part Number	HMCG88AGBRA190N
	Size (GB)	64
	Channel	2
BIOS	Vendor	Intel Corporation
	Version	KVLDCRB1.KWO.0030.D39.2502281359
	Build Date	02/28.2025
Network Adapter 1	Device	Intel® Ethernet Network Adapter E830-XXVDA2
	Link Speed	25Gbps
	Firmware	1.00 0x80016fcb 1.3832.0
	Driver	ice 2.3.3
Network Adapter 2	Device	NVIDIA MCX713104AS-ADAT ConnectX-7 HHHL
	Link Speed	25Gbps
	Firmware	28.37.1014 (MT_0000000849)
	Driver	mlx5_core
Network Adapter 3	Device	Broadcom BCM957504-P425G
	Link Speed	25Gbps
	Firmware	226.0.145.1/pkg 226.1.107.1
	Driver	bnxt_en

Software		
OS	Vendor	Ubuntu 22.04 LTS
	Kernel Version	6.9
	Kernel boot parameters	ro quiet splash split_lock_detect=off vt.handoff=7
	OS Tunings	tuna -cpus=1-3 -isolate echo -1 > /proc/sys/kernel/sched_rt_runtime_us
Load Software		iPerf 2.1.5 pthreads
Time Software		LinuxPTP 4.4

Load Configuration

For load testing, multiple iPerf flows were used to generate system contention. One flow transmits UDP packets at the maximum iPerf rate from the device under test to the Calnex GM. The other flow uses a second network adapter with two ports in separate network namespaces connected back-to-back to generate bidirectional 25 Gbps TCP streams between the ports.

Important Information About this Report

CONTRIBUTORS

Russ Fellows

VP, Labs | Signal65

Cameron Moccari

Operations and Project Management Director | Signal65

PUBLISHER

Ryan Shrout

President and GM | Signal65

INQUIRIES

Contact us if you would like to discuss this report and Signal65 will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Signal65." Non-press and non-analysts must receive prior written permission by Signal65 for any citations.

LICENSING

This document, including any supporting materials, is owned by Signal65. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of Signal65.

DISCLOSURES

Signal65 provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT SIGNAL65

Signal65 is a leading research organization specializing in enterprise AI infrastructure optimization and deployment strategies. Our lab focuses on evaluating and optimizing AI hardware and software solutions for real-world enterprise applications, with particular expertise in large language models, retrieval-augmented generation systems, and distributed AI architectures.

For more information, visit signal65.com or contact research@signal65.com



IN PARTNERSHIP WITH



CONTACT INFORMATION

Signal65 | signal65.com