



AMD LOW LATENCY SOLUTIONS

AMD EPYC™ 9005 Series Processors and AMD Alveo™ X3 Adapter

WHITEPAPER | 2025

***Authors:** Ashwin Thiagarajan, Sr.Manager; Paul Emberson, PMTS; Nimisha Raut, SMTS*

***Credits:** Faisal Dada, Director; Robert Winter, Director; Marcus Yoon, Sr.Manager*

Innovation in Ethernet network adapter technology has advanced to address a wide range of use cases, applications, and services required by diverse markets. The financial technology (or **Fintech**) market requires low latency, wire speed packet performance, deterministic packet jitter, and specialized multicast processing. These requirements are addressed with AMD Solarflare™ OpenOnload™ and Network adapters, optimization of operating system and CPUs, BIOS and host software. This paper describes how AMD EPYC™ CPUs and Alveo™ X3522 adapters, paired with OpenOnload, deliver ultra-low latency and deterministic jitter performance for fintech workload.

- **Network latency** is the time spent by a packet to traverse from transmitter to receiver, reported in nano seconds.
- **Network jitter and jitter determinism** measures the variation of packet latency from mean latency, reported typically in 90th percentile or above.

FINTECH REQUIREMENTS

AMD provided Fintech solutions:

- Performant and high core count CPUs (AMD EPYC™ and Ryzen™ CPUs)
- Specialized low latency network adapters (AMD Solarflare), FPGAs and AMD Alveo Adapters
- Kernel bypass protocol acceleration software (AMD Solarflare OpenOnload)

AMD BETTER TOGETHER; WORKING TO PROVIDE LOWEST LATENCY

Historically, competitor x86_64 CPU solutions with their monolithic CPU architecture, provided better system and network latencies. With lower networking latency in mind, AMD server and networking teams have been working together to optimize the networking latency and integration of AMD Solarflare adapter with AMD Ryzen and EPYC CPUs. With advanced optimization, customers can benefit through enhanced performance, low power consumption, high core count, and low Total Cost of Ownership (TCO).


X3522 Network Latency Comparison

Latency performance is measured with **eflatency** and **sfnt-pingpong** benchmarking tools. Eflatency is a network layer L2 aware single packet, multi-iteration, single filter application designed to measure round trip delay between a sender and receiver. Sfnt-pingpong is a network layer L3 aware application which can handle multiple packets, and has support for multiple filters, designed to measure latency between a sender and receiver.


X3522: Performance - Eflatency

The following chart compares EPYC 9375F and Ryzen 7950X for latencies. With 9375F, while an (+25ns) latency was observed for 4B payload size when compared with 7950X, comparable results (+11ns) are observed for payload sizes ranging from 0B to 648B with Alveo X3 adapter.

X3: 10G: Eflatency - AMD CPUs

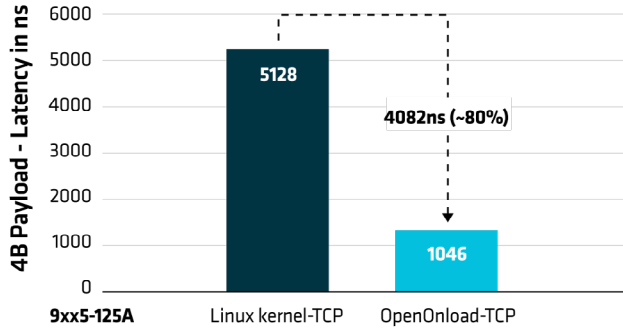


X3: 10G: Latency Performance for 4B



X3522: Performance – Onload

Turin™ 9375F CPU: X3-Series, 10G Linux® kernel-TCP vs OpenOnload™ - TCP latency

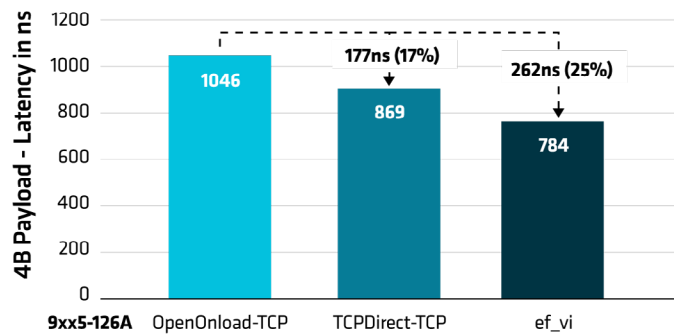


Onload TCP: An average 80% reduction in latency for 4B payload when compared to kernel

With Onload acceleration, sfnt-pingpong TCP test with 4B payload size, produced an average 80% reduction in latency compared to the same application executed with kernel network stack.

For 0B to 648B, the same tests produced 76% reduction in latency when average value for the range is compared.

Turin™ 9375F CPU: X3, 10G Onload™ Family latency

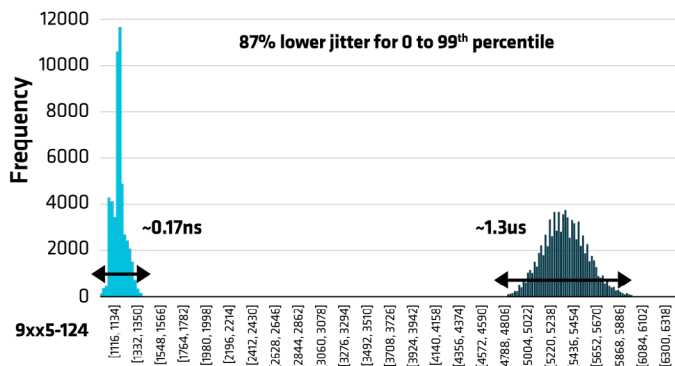


An average 17% reduced latency with TCPDirect & ~25% reduced latency with ef_vi for 4B payload

Onload offers 3 APIs for reducing latency for diverse customer use cases. At the expense of development effort, compared to zero touch Onload, customers can develop applications with additional latency reduction.

For 0B to 648B, the same tests produced an average latency reduction of ~12.5% with TCPDirect and ~19.2% with ef_vi APIs.

100K Loop Determinism: kernel-TCP vs Onload TCP



An average 87% Lower Jitter for deterministic performance of 256B for 100K iterations

With OpenOnload, 87% lower Jitter and lower deterministic jitter performance was achieved, compared to the same application on the kernel TCP/IP stack.

The more constrained network jitter provided by OpenOnload allows applications more deterministic network access. This in turn offers fairness in trading, deterministic and faster trade execution.

AMD FINTECH LATENCY REDUCTION – AN INTRODUCTION

AMD has taken an end-to-end approach of helping in reducing overall latency in real customer use cases.

Using Onload, the AMD Solarflare kernel bypass technology, along with the implementation of Cut Through Programmable IO (CTPIO) in the NIC, addresses transmit latency. On the receive side, by implementing the polling and using hardware filters configured by the software, the receive Direct Memory Access (DMA) transactions are routed to specific queues managed by the Onload. After the DMA, NIC will generate an event to notify the host that the packet delivery has been completed. An application pinned on a specific CPU will then consume the data.

This holistic approach provides best end to end latency by accelerating Ethernet frames on both transmit and receive.

AMD FINTECH LATENCY REDUCTION – A NOVEL APPROACH

Achieving a low latency is an art, a scientific approach, starting from bare metal hardware to system bring up, network bring up and finally placing the infrastructure in an optimal location. This document provides overall guidance on various resources available for customers when planning for low latency with EPYC 9005 and AMD NICs.

[AMD EPYC™ 9005 BIOS & Workload Tuning guide](#)

[AMD Solarflare X3522 NIC resources](#)

[Low latency tuning for AMD EPYC powered servers](#)

[OpenOnload resources](#)

NEXT STEPS

[Low Latency Tuning For AMD EPYC CPU Powered Servers](#)

[Release Notes](#)

Onload and associated debug and test tools are publicly available for evaluation with X2 or X3 adapters. Please contact **[AMD sales](#)** for more information.

The Ryzen AMD processors are publicly available and can be purchased through AMD authorized distributors and channel partners. Please contact **[AMD sales](#)** for more information. The 5th generation EPYC based servers are available from OEMs.

END NOTES

9xx5-125A: Low latency – OpenOnload vs kernel latency: AMD Internal Testing as of 2/20/2025. Workload config: SFNT-pingpong tool ran with Onload 8.1.3.40 vs the same application with kernel TCP-IP stack. Latency in ns for payload lengths ranging from 4-648B. System Config: 2x1P AMD EPYC 9375F powered reference system (32C, 32GB DDR5-4800, RHEL 9.5 5.14.0-503.14.1.el9_5.x86_64, BIOS RPUT1003B, OS and BIOS at default settings) with AMD Alveo™ X3522 (fw.bundle_id 1.28.127.1158); An average of 76% lower latency when using OpenOnload-TCP vs kernel TCP-IP stack across payload size 4B to 648B; 4B payload result v- avg 1046ns for OpenOnload-TCP vs avg 5128ns for kernel; or avg 79.6% reduction in latency; Results may vary due to factors including system configurations, software versions, and BIOS settings.

9xx5-126A: Onload family provides flexible APIs: AMD Internal Testing as of 2/20/2025. Workload config: SFNT-pingpong tool ran with Onload 8.1.3.40, eflatency (ef_vi), TCPDirect-tcp, and kernel TCP-IP stack. Latency in ns for payload lengths ranging from 4-648B. System Config: 2x1P AMD EPYC 9375F powered reference system (32C, 32 GB DDR5-4800, RHEL 9.5 5.14.0-503.14.1.el9_5.x86_64, BIOS RPUT1003B, OS and BIOS at default settings) with AMD Alveo X3522 (fw.bundle_id 1.28.127.1158) ; 4B payload result vs 1046ns for OpenOnload-TCP: - 869ns for TCPdirect-TCP (~16.9% reduction in latency) ; - 784ns for Eflatency (~25.0% reduction in latency) ; Payload average of 4 to 648B: ~12.5% TCPdirect-TCP latency reduction ;~19.2% Eflatency latency reduction; Results may vary due to factors including system configurations, software versions, and BIOS settings.

9xx5-124: Determinism – OpenOnload Jitter Performance: AMD Internal Testing as of 2/20/2025. Workload config: SFNT-pingpong tool ran with Onload 8.1.3.40 vs kernel TCP/IP Stack . Measured 256B payload latency with 100K iterations for deterministic low jitter performance. System Config: AMD EPYC 9375F powered production system (32C, 32GB DDR5-4800, RHEL 9.5 5.14.0-503.14.1.el9_5.x86_64, BIOS RPUT1003B) with AMD Alveo™ X3522 (fw.bundle_id 1.28.127.1158) - Min spread for 0-99th percentile of 1382 ns (kernel-TCP) - Min spread for 0-99th percentile of 175 ns (Onload-TCP) - For 87% (7x) lower jitter Results may vary due to factors including system configurations, software versions, and BIOS settings

9xx5-127A: Latency performance: AMD Internal Testing as of 03/05/2025 (EPYC systems) and 6/14/2024 (Ryzen system only) ; Workload config: OpenOnload™/eflatency Version 8.1.3.40, Latency in ns for payload lengths ranging from 4B and 0-648B. ; System Config: 2x1P AMD EPYC 9375F powered reference system (32C, 32GB DDR5-4800, RHEL 9.5 Kernel: 5.14.0-503.14.1.el9_5.x86_64, BIOS RPUT1003B, NPS4, SMT=off, Power Determinism, Power states) with AMD Alveo™ X3522 (fw.bundle_id 1.28.127.1158); System Config: AMD Ryzen 9 7950X powered SuperMicro H13SAE-MF system (16C, 32GB DDR5-4800, RHEL 9.4 Kernel 5.14.0-362.18.1, BIOS 1.0a, SMT=off) with AMD Alveo™ X3522 (fw.bundle_id 1.28.127.1158); 4B payload result: 9375F to 7950X: 24.5 ns higher latency compared to 7950X; Results may vary due to factors including system configurations, software versions, and BIOS settings

DISCLAIMERS

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED "AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

COPYRIGHT NOTICE

©2025 Advanced Micro Devices, Inc. All Rights Reserved. AMD, the AMD arrow logo, EPYC, Ryzen, Alveo, Solarflare, Xilinx, Xilinx logo, Onload, OpenOnload, EnterpriseOnload, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this presentation are for identification purposes only and may be trademarks of their respective companies.