

May 7, 2024



Optiver Chooses AMD Enterprise Portfolio to Power its Data Center Modernization, Enabling New Era of Compute and AI

Optiver utilizes high-performance AMD solutions including EPYC CPUs, Solarflare Ethernet adapters, Virtex FPGAs and Alveo Accelerators

SANTA CLARA, Calif., May 07, 2024 (GLOBE NEWSWIRE) -- Today, [AMD](#) (NASDAQ: AMD) announced that [Optiver](#) – a leading global market maker trading on 100+ exchanges – is using a broad range of AMD high-performance compute engines to help further its mission of improving financial markets by building a modern infrastructure that serves as the backbone of the organization. Using AMD EPYC™ processors, AMD Solarflare low-latency Ethernet adapters, AMD Virtex™ FPGAs, and AMD Alveo™ adaptable accelerators, Optiver is addressing a variety of technical challenges in capital markets to build systems that make financial markets more advanced.

"We're a team of researchers and engineers driven by a common vision: building systems that make financial markets better," said Alexander Itkin, U.S. chief technology officer, Optiver. "We are constantly pushing the limits, whether it's challenging the speed of light to minimize latency or building world-leading computational scale. We need computing infrastructure that enables us to navigate the complexities of the capital markets in nanoseconds. For us, this is only possible with high-performance products from AMD, which help us successfully and continuously adapt to financial markets."

"AMD is uniquely positioned to meet the needs of commercial customers with a robust and diverse portfolio of compute engines that enable impressive performance and energy efficiency; two critical components of the modern data center," said Archana Vemulapalli, corporate vice president, Commercial Enterprise, AMD. "Our work with Optiver is a great example of how we work together with customers to push the limits of what is possible, while demonstrating the capabilities of an AMD commercial solution and how it can address both present-day and future challenges for our customers."

The specific AMD enterprise products that enable Optiver to modernize its infrastructure and unlock speed and efficiency for its business include:

- **AMD EPYC CPUs:** Optiver is using [4th Gen AMD EPYC processors](#) to build out large-scale compute needed to drive data analysis and deliver core density for the company's data center modernization and workload consolidation efforts.
- **Solarflare Ethernet Adapters:** Optiver also uses the [AMD Solarflare X2 series Ethernet network adapters](#) designed for high-performance electronic trading environments. These adapters provide low-level control of network interfaces at the application level, allowing Optiver to better design systems to scale and stay competitive.

- **Virtex UltraScale+ FPGAs:** The exponential growth of data demands smart network and data center solutions that reach beyond fixed-function silicon devices to deliver maximum throughput, high data processing capability, and the flexibility to adapt to evolving connectivity standards. The [AMD Virtex™ UltraScale+™ VU23P FPGA](#) helps Optiver accelerate its systems while remaining endlessly adaptable.
- **Alveo Accelerators:** Optiver utilizes the [AMD Alveo U55C](#) and [Alveo UL3524](#) adaptable accelerator cards, the latter of which delivers a 7X latency improvement over prior generation FPGA technology¹.

“Our ultimate goal was to get faster access to data and information that is critical to meet the unique needs and workflows that the financial industry demands,” said Kevin Sprague, U.S. hardware lead, Optiver. “The high-performance Alveo ultra-low latency cards are critical to helping us move more information, faster. Our collaboration with AMD helped us develop the products and features our industry demands. Additionally, our joint work has allowed us to expand and optimize our offering by leveraging not only AMD FPGAs, but also AMD EPYC CPUs to build out high-performance trading systems and infrastructure.”

AMD helps enterprise customers address the challenges of today with a robust portfolio of innovative and proven solutions. Optiver illustrates the ability of the leadership AMD compute portfolio, where individual devices in tandem form intelligent solutions greater than the sum of their parts, delivering unparalleled value to companies and customers.

Supporting Resources

- Follow AMD on [LinkedIn](#)
- Follow AMD on [X](#)

About AMD

For more than 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies. Billions of people, leading Fortune 500 businesses, and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) [website](#), [blog](#), [LinkedIn](#), and [X](#) pages.

©2024 Advanced Micro Devices, Inc. All rights reserved. AMD, Alveo, EPYC, Solarflare, Virtex and combinations thereof are trademarks of Advanced Micro Devices, Inc. Arm is a registered trademark of Arm Limited (or its subsidiaries or affiliates) in the US and/or elsewhere. Other names used herein are for informational purposes only and may be trademarks of their respective owners.

¹ Testing conducted by AMD Performance Labs as of 8/16/23 on the Alveo UL3524 accelerator card, using Vivado™ Design Suite 2023.1 and running on Vivado Lab (Hardware Manager) 2023.1. Based on the GTF Latency Benchmark Design configured to enable GTF transceivers in internal near-end loopback mode. GTF TX and RX clocks operate at same frequency of ~644MHz with a 180 degrees phase shift. GTF Latency Benchmark Design measures latency in hardware by latching value of a single free running counter. Latency is measured as the difference between when TX data is latched at the GTF transceiver and when TX data is latched at the GTF receiver prior to routing back into the FPGA fabric. Latency measurement does not include protocol overhead, protocol framing, programmable logic (PL) latency, TX PL interface setup time, RX PL interface clock-to-out, package flight

time, and other sources of latency. Benchmark test was run 1,000 times with 250 frames per test. Cited measurement result is based on GTF transceiver “RAW Mode”, where PCS (physical medium attachment) of the transceiver passes data ‘as-is’ to FPGA fabric. Latency measurement is consistent across all test runs for this configuration. System manufacturers may vary configurations, yielding different results. ALV-10

Contact:

Aaron Grabein

AMD Communications

(512) 602-8950

Aaron.grabein@amd.com

Suresh Bhaskaran

AMD Investor Relations

+1 408-749-2845

Suresh.Bhaskaran@amd.com



Source: Advanced Micro Devices, Inc.