

September 27, 2023



AMD Unveils Purpose-Built, FPGA-Based Accelerator for Ultra-Low Latency Electronic Trading

— New AMD Alveo fintech accelerator card provides trading firms and brokerages with breakthrough trade execution performance at nanosecond speed and AI-enabled trading strategies —

— Solution partners Alpha Data, Exegy and Hypertec add to growing ecosystem of ultra-low latency solutions for fintech market —

SANTA CLARA, Calif., Sept. 27, 2023 (GLOBE NEWSWIRE) -- [AMD](#) (NASDAQ: AMD) today announced the [AMD Alveo™ UL3524 accelerator card](#), a new fintech accelerator designed for ultra-low latency electronic trading applications. Already deployed by leading trading firms and enabling multiple solution partner offerings, the Alveo UL3524 provides proprietary traders, market makers, hedge funds, brokerages, and exchanges with a state-of-the-art FPGA platform for electronic trading at nanosecond (ns) speed.

The Alveo UL3524 delivers a 7X latency improvement over prior generation FPGA technology¹, achieving less than 3ns FPGA transceiver latency² for accelerated trade execution. Powered by a custom 16nm Virtex™ UltraScale™+ FPGA, it features a novel transceiver architecture with hardened, optimized network connectivity cores to achieve breakthrough performance. By combining hardware flexibility with ultra-low latency networking on a production platform, the Alveo UL3524 enables faster design closure and deployment compared to traditional FPGA alternatives.

“In ultra-low latency trading, a nanosecond can determine the difference between a profitable or losing trade,” said Hamid Salehi, director of product marketing at AMD. “The Alveo UL3524 accelerator card is powered by the lowest latency FPGA transceiver from AMD—purpose-built to give our fintech customers an unprecedented competitive advantage in financial markets.”

Hardware Flexibility and AI-Enabled Trading Strategies

Featuring 64 ultra-low latency transceivers, 780K LUTs of FPGA fabric, and 1,680 DSP slices of compute, the Alveo UL3524 is built to accelerate custom trading algorithms in hardware, where traders can tailor their design to evolving strategies and market conditions. Supported by traditional FPGA flows using Vivado™ Design Suite, the product comes with a suite of reference designs and performance benchmarks that allow FPGA designers to quickly explore key metrics and develop custom trading strategies to specification, backed by global support from AMD domain experts.

To simplify the increasing adoption of AI in the algorithmic trading market, AMD is providing developers with the open-sourced and community-supported [FINN development framework](#). By using PyTorch and neural network quantization techniques, the FINN project enables

developers to reduce the size of the AI models while retaining accuracy, compiling to hardware IP, and integrating the network model into the algorithm's datapath for low latency performance. As an open-source initiative, the solution gives developers flexibility and accessibility to the latest advancements as the projects evolve.

Enabling a Growing Ecosystem of Ultra-Low Latency Fintech Solutions

The Alveo UL3524 and purpose-built FPGA technology are enabling strategic partners to build custom solutions and infrastructure for the fintech market. Currently available partner solutions include offerings from Alpha Data, Exegy, and Hypertec.

The AMD Virtex™ UltraScale+ VU2P FPGA powering the Alveo UL3524 accelerator card is enabling ultra-low latency appliances from Alpha Data.

"The new Virtex UltraScale+ FPGA from AMD brings a step change to ultra-low latency trading and networking," said David Miller, managing director of Alpha Data. "We've developed the ADA-R9100 rack-mount appliance which allows customers to easily access the full potential of the new AMD FPGA device."

Exegy, a provider of end-to-end, front-office trading solutions, is supporting the Alveo UL3524 card with its nxFramework, a software and hardware development environment tailored for creating and maintaining ultra-low latency FPGA applications within the financial industry.

"By combining the pioneering ultra-low latency FPGA technology from AMD with Exegy's expertise in capital markets, we're able to deliver a comprehensive solution that addresses the ever-increasing optimization needed to build the trading infrastructure of tomorrow," said Olivier Cousin, director of FPGA solutions at Exegy.

Hypertec has optimized its ORION HF X410R-G6 High Frequency Server for the Alveo UL3524 with a custom cooling system to deploy in a 1U server form factor.

"The engineers at Hypertec specifically designed the HF X410R-G6 to extract the best out of the capabilities and speed of the Alveo UL3524 platform, catering our solution to the most demanding low-latency tasks," said David Lim, director of product marketing, Hypertec.

The AMD Alveo UL3524 accelerator card is currently in production and shipping to global financial services customers.

Supporting Resources

- Learn more about the [Alveo UL3524 accelerator card](#)
- Become a fan of AMD on [Facebook](#)
- Follow AMD on [Twitter](#)

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 [Twitter](#) pages.

AMD, the AMD Arrow logo, Alveo, Virtex, Vivado and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other names 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

² Based on simulation comparison between Virtex UltraScale+ GTY transceivers and ultra-low latency GTF transceivers.

Contact:

Mike Sanchez
AMD Communications
(209) 262-7458
m.sanchez@amd.com

Suresh Bhaskaran
AMD Investor Relations
(408) 749-2845
Suresh.bhaskaran@amd.com

A photo accompanying this announcement is available at
<https://www.globenewswire.com/NewsRoom/AttachmentNg/b51a1b28-7ee7-45b5-9028-e18f6de2634f>.



Source: Advanced Micro Devices,
Inc.

AMD Alveo UL3524 Accelerator Card



Image of AMD Alveo UL3524 Accelerator Card - angled view