

April 6, 2023



AMD Launches First 5nm ASIC-based Media Accelerator Card to Power New Era of Interactive Media Services at Scale

Purpose-built video processing architecture featuring an AV1 accelerated pipeline delivers 32x 1080p streams per card with AI-enabled video quality optimization

SANTA CLARA, Calif., April 06, 2023 (GLOBE NEWSWIRE) -- [AMD](#) (NASDAQ: AMD) today announced the AMD Alveo™ MA35D media accelerator featuring two 5nm, ASIC-based video processing units (VPUs) supporting the AV1 compression standard and purpose-built to power a new era of live interactive streaming services at scale. With over 70% of the global video market being dominated by live content¹, a new class of low-latency, high-volume interactive streaming applications are emerging such as watch parties, live shopping, online auctions, and social streaming.

The Alveo MA35D media accelerator delivers the high channel density, with up to 32x 1080p60 streams per card, power efficiency and ultra-low-latency performance critical to reducing the skyrocketing infrastructure costs now required for scaling such compute intensive content delivery. Compared to the previous generation Alveo U30 media accelerator, the Alveo MA35D delivers up to 4x higher channel density², 4x max lower latency in 4K³ and 1.8x greater compression efficiency⁴ to achieve the same VMAF score – a common video quality metric.

“We worked closely with our customers and partners to understand not just their technical requirements, but their infrastructure challenges in deploying high-volume, interactive streaming services profitably,” said Dan Gibbons, general manager of AECG Data Center Group, AMD. “We developed the Alveo MA35D with an ASIC architecture tailored to meet the bespoke needs of these providers to reduce both capital and operating expenses for delivering immersive experiences to their users and content creators at scale.”

Purpose-Built Video Processing Unit

The Alveo MA35D utilizes a purpose-built VPU to accelerate the entire video pipeline. By performing all video processing functions on the VPU, data movement between the CPU and accelerator is minimized, reducing overall latency and maximizing channel density with up to 32x 1080p60, 8x 4Kp60, or 4x 8Kp30 streams per card. The platform provides ultra-low latency support for the mainstream H.264 and H.265 codecs and features next-generation AV1 transcoder engines delivering up to a 52% reduction in bitrate for bandwidth savings versus a comparable software implementation⁵.

“AMD’s announcement of the new Alveo MA35D add-in card is an exciting advancement of video acceleration for data centers and is an important step in building out a fully-fledged ecosystem to support royalty-free, high-definition video devices, products, and services,”

said Matt Frost, Alliance for Open Media Chair. “Live streaming providers are looking for higher density, lower power, lower latency AV1 solutions and by addressing these, Alliance members such as AMD are helping facilitate AV1 deployment and overall adoption.”

AI-Enabled, Intelligent Video Pipeline

The accelerator features an integrated AI processor and dedicated video quality engines designed to improve the quality of experience at reduced bandwidth. The AI processor evaluates content, frame-by-frame, and dynamically adjusts encoder settings to improve perceived visual quality while minimizing bitrate. Optimization techniques include region-of-interest (ROI) encoding for text and face resolution, artifact detection to correct scenes with high levels of motion and complexity, and content-aware encoding for predictive insights for bitrate optimization.

Cost-Effectively Scale Interactive Media

Scaling high-volume streaming services requires maximizing the number of channels per server while minimizing power and bandwidth-per-stream. By delivering up to 32x 1080p60 streams per card at 1 watt per stream⁶, a 1U rack server equipped with 8 cards delivers up to 256 channels to maximize the number of streams per server, rack or data center.

Software Dev Kit and Product Availability

The platform is accessible with the AMD Media Acceleration software development kit (SDK), supporting the widely used FFmpeg and Gstreamer video frameworks for ease of development.

Alveo MA35D media accelerators are sampling now with production shipments expected in Q3. To accelerate development, an Early Access Program is available to qualified customers with comprehensive documentation and software tools for architectural exploration.

Supporting Resources

- Learn more about the [Alveo MA35D Media Accelerator](#)
- Check out the Alveo MA35D in the AMD booth at NAB, North Hall N2158
- Learn more about the [Alliance for Open Media](#) and its member companies
- 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.

Contact:

Mike Sanchez

AMD Communications

(209) 262-7458

m.sanchez@amd.com

¹ [Source: Bluewave Consulting and Research, March 2022](#)

² In published specifications, the Alveo MA35D supports up to 32 1080p60 streams, while the Alveo U30 supports up to 8. Channel density ratios remain the same regardless of resolution. ALV-002

³ In published specifications, the Alveo MA35D delivers 4X lower latency at 8ms vs. Alveo U30 delivering 4K H.264 at 32ms, based on lowest latency capability of each platform. ALV-005

⁴ Based on testing by AMD Labs in April 2023, using the VMAF scores of a Alveo MA35D AV1 encode compared to Alveo U30 H.264 encode across (13) publicly available video files at various resolutions and bitrates. Actual results may vary. ALV-009

⁵ Based on testing by AMD Labs in March 2023, using the VMAF scores of Alveo MA35D H.264 encode, H.265 encode, and AV1 encode compared to the VMAF score of an open source x264 veryfast SW model across (13) publicly available video files at various resolutions and bitrates. Actual results may vary. ALV-006

⁶ Typical power for 8 4K streams or 32 1080p60 streams estimated at 35W, based on preliminary testing and subject to change. 50W Total Thermal Design Power (TDP)



Source: Advanced Micro Devices, Inc.