

SOFTWARE-DEFINED BROADCAST AV WITH AMD RYZEN™ AND EPYC™ EMBEDDED PROCESSORS



together we advance_AV

OVERVIEW

The broadcast and professional AV industry is undergoing a fundamental shift: purpose-built hardware appliances are giving way to software-defined architectures that deliver greater flexibility, faster deployment and cloud-style operating models, without sacrificing the deterministic performance that live video demands.

Software-defined workflows now span encoding, decoding, playout, IP routing, media server operation and virtualized production. The processor becomes a key architectural decision as CPU capability directly affects channel density, video quality, power consumption and total cost of ownership.

AMD Ryzen™ Embedded and Ryzen AI Embedded processors combine x86 compute, integrated GPU and AMD XDNA™ NPU in compact form factors for broadcast encoders, production switchers and AV-over-IP routing appliances.

For server-class workloads, AMD EPYC™ Embedded processors deliver industry-leading performance-per-watt and I/O density for multi-channel encoding and virtualized production.

THE SOFTWARE-DEFINED ARCHITECTURE ADVANTAGE

AMD Ryzen Embedded processors replace fixed-function hardware with flexible x86 platforms updated and scaled through software alone. Integrated AMD Radeon™ GPU graphics accelerate real-time encode/decode (H.264, H.265, AV1). The integrated AMD XDNA NPU reduces core utilization and accelerates AI-enhanced applications such as automated camera switching, content tagging and scene analysis, with full-width AVX-512.

This software-defined approach decouples capability from hardware cycles. Broadcasters and AV integrators can deploy new codecs, add AI features or reconfigure channels without replacing equipment, reducing capital expenditure and operational risk.

With 10+ year industrial longevity and broad OS and hypervisor support, AMD Ryzen Embedded processors provide a stable, future-ready foundation from edge appliances to control rooms.

SOFTWARE-DEFINED FLEXIBILITY

Deploy, update and reconfigure AV workflows entirely through software. x86 software compatibility enables straightforward migration from existing platforms.

INTEGRATED AI AND MEDIA ACCELERATION

Leverage built-in GPU and NPU for real-time video and AI processing on a single device.

SCALABLE FROM EDGE TO CLOUD

AMD Ryzen Embedded processors for power-efficient edge appliances; AMD EPYC Embedded processors for dense server workloads with lower power and cooling costs.

KEY BENEFITS

MULTI-CHANNEL ENCODING

- Run multiple 4K encode/decode streams with integrated GPU and CPU processing on a single platform.

BUILT-IN AI PROCESSING


- AMD XDNA NPU enables real-time AI inference at the edge without external accelerators.

10+ YEAR PLATFORM LONGEVITY


- Industrial-grade embedded processors with extended lifecycle support.

CONVERGED AV WORKFLOWS

- Run encoding, switching, playout and AI on one x86 platform with full OS and software ecosystem support.




SOFTWARE-DEFINED ENCODING & TRANSCODING




Flexible x86 platforms adapt to new codecs through software alone. AMD Ryzen Embedded processors offer multi-channel 4K encoding via an integrated AMD Radeon GPU, consistently outperforming competing integrated graphics across Vulkan® and OpenGL® workloads.^{1,2} The AMD Advanced Media Framework (AMF) SDK unlocks hardware-accelerated encode/decode pipelines for high channel density in compact, power-efficient systems.

AI-POWERED LIVE PRODUCTION




AI is transforming live production by automating camera selection, graphic insertion and real-time content tagging. AMD Ryzen AI Embedded processors integrate a dedicated AMD XDNA NPU alongside the CPU and GPU, running on-device inference for detection, recognition and classification, with unified memory architecture that dynamically shares bandwidth across concurrent AI and media workloads. Compact systems run AI-assisted directing, automated highlight clipping and metadata generation at the edge, reducing crew requirements and accelerating content turnaround.

AV-OVER-IP MEDIA ROUTING



AV-over-IP protocols such as ST 2110, NDI, Dante AV and IPMX are replacing SDI matrices with software-based routing that runs on IP network infrastructure. AMD Ryzen Embedded processors provide the multi-threaded CPU performance, integrated GPU, integrated 10 Gb Ethernet and native AV1 hardware acceleration to encode, decode and route multiple AV-over-IP streams. Software-defined routers scale from small studios to large venue installations and can be reconfigured or expanded without replacing hardware.

VIRTUALIZED PRODUCTION CONTROL ROOMS



Virtualized control rooms replace dedicated equipment racks with shared compute pools that dynamically allocate encoding, switching and graphics. AMD EPYC Embedded processors deliver the core density, memory bandwidth and up to 128 PCIe® Gen5 lanes needed to run multiple production workloads concurrently in virtualized or containerized environments. Scale up for live events then reallocate resources when production ends, eliminating idle hardware.

THE AMD ADVANTAGE

AMD RYZEN AI EMBEDDED X100 VS. INTEL® CORE™ ULTRA SERIES 3	
HIGH-PERFORMANCE COMPUTE	- AMD Ryzen AI Embedded X100 fields all-classic-core compute. Intel Core Ultra Series 3 headlines 16 cores but exposes just 4 performance cores (4P+8E+4LP-E), throttling sustained AV workloads. AMD delivers up to 16 classic cores/32 threads and an estimated 1.5x higher multi-thread / 1.1x higher single-thread SPECrate®2017_int_base throughput with no efficiency-core penalty.³ - Under load, X100 sustains 1.3x higher memory bandwidth (STREAM)⁴ than Intel Core Ultra Series 3, keeping encode and multiview pipelines fed on up to 273 GB/s bandwidth and a 32 MB MALL cache.
INTEGRATED GPU MEDIA ENGINES	Integrated AMD Radeon GPU encodes/decodes H.264, H.265, and AV1 for multi-channel 4K and real-time multiview with no discrete GPU. Most Intel Core Ultra Series 3 SKUs ship only a small 4x Xe3 graphics tile, while X100 is expected to deliver an estimated 1.4x higher Vulkan¹ and 1.7x higher OpenGL² graphics for superior rendering and compute acceleration.
ON-DEVICE AI PROCESSING	AMD XDNA NPU accelerates real-time inference – object detection, scene analysis, automated production – with no external accelerators. While Intel Core Ultra Series 3 is capped at a 256-bit AVX2 path, AMD full-width AVX-512 provides twice the vector width. For generative AI, X100 delivers 3.5x higher token generation and 1.4x faster time-to-first-token than Intel Core Ultra X7 358H.⁵
HYBRID EDGE-CLOUD PROCESSING	On-premise and remote production via compute partitioning over PCIe Gen5, USB4 and 2x 10 GbE TSN networking – 4x the per-port bandwidth of Intel Core Ultra Series 3's 3x 2.5 GbE – moving AV-over-IP, NVMe and accelerators without external switches.
PLATFORM SECURITY	AMD Memory Guard and Secure Boot with hardware-rooted protection keep AV-over-IP streams and production resilient – embedded security beyond Intel Core Ultra Series 3's client-focused device.
10+ YEAR PLATFORM LONGEVITY	Industrial-grade parts on mature, monolithic TSMC 4 nm silicon – one node, one fab – with multi-gen socket compatibility, a -40°C to +105°C junction temp and 10-year planned availability rated for 24/7 operation, unlike Intel Core Ultra Series 3, which spans three nodes, two fabs and proprietary packaging.

NEXT STEPS

- For more information on AMD x86 Embedded processors, visit www.amd.com/en/products/embedded.html
- To learn more about the AMD Advanced Media Framework, visit gpuopen.com/advanced-media-framework
- For more on Broadcast AV solutions, visit www.amd.com/broadcastav
- To learn more about AMD X100 advantages, visit www.amd.com/en/products/embedded/ryzen-ai/x100-advantage.html

ENDNOTES

1. Measurements and projections conducted by AMD as of May 2026 on an AMD Ryzen AI Max+ 395 processor (OPN 100-000002199) configured to reflect Ryzen AI Embedded X199 specifications (AMD Maple CRB, 5.1 GHz CPU, 2.9 GHz GPU, with sustained operation at 45W TDP/sPPT/fPPT, STAPM disabled and 64 GB soldered LPDDR5X-8000), compared to an Intel Core Ultra X7 358H (MSI Prestige 16 Flip AI+ C3MTG MS-2622, 5.1 GHz CPU, 2.5 GHz GPU, PL1 = 45W, PL2 = 64W, MSI BIOS E2622IMS.10E, 32 GB LPDDR5X-8533, measured 30W OEM Sustained Power Limit). Results reflect the relative geometric mean of offscreen Vulkan benchmark frame rates from GFXBench 5.0.0 on Ubuntu 24.04, comparing measured AMD performance at 45W sustained power to 45W Intel performance projected from 30W sustained power measurements, using scaling factors derived from public benchmark data for the 358H. System manufacturers may vary configurations, yielding different results. Results may vary based on configuration, settings, usage, and other factors. (REX-002)
2. Measurements and projections conducted by AMD as of May 2026 on an AMD Ryzen AI Max+ 395 processor (OPN 100-000002199) configured to reflect Ryzen AI Embedded X199 specifications (AMD Maple CRB, 5.1 GHz CPU, 2.9 GHz GPU, with sustained operation at 45W TDP/sPPT/fPPT, STAPM disabled and 64 GB soldered LPDDR5X-8000), compared to an Intel Core Ultra X7 358H (MSI Prestige 16 Flip AI+ C3MTG MS-2622, 5.1 GHz CPU, 2.5 GHz GPU, PL1 = 45W, PL2 = 64W, MSI BIOS E2622IMS.10E, 32 GB LPDDR5X-8533, measured 30W OEM Sustained Power Limit). Results reflect the relative geometric mean of offscreen OpenGL frame rates from GFXBench 5.0.0 on Ubuntu 24.04, comparing measured AMD performance at 45W sustained power to 45W Intel performance projected from 30W sustained power measurements, using scaling factors derived from public benchmark data for the 358H. System manufacturers may vary configurations, yielding different results. Results may vary based on configuration, settings, usage, and other factors. (REX-003)
3. Estimated by AMD as of May 2026 on the AMD Ryzen AI Max+ 395 (OPN 100-000002199) at 45W SPL (sPPT=45W, fPPT=64W) and the Intel Core Ultra X7 358H at 30W SPL (PL1=45W, PL2=64W) using SPEC CPU2017_int_base (configuration: GCC-15, rate, harness cpu2017-1.1.9, int base sub score) on test systems comprising:- Ryzen AI Max+ 395 (OPN 100-000002199): AMD Reference Maple Motherboard, 64 GB LPDDR5X-8000 SK Hynix RAM, Kingston Technology KV3000 NVMe SSD 512 GB PCIe Gen4 512 GB storage, Linux Ubuntu 24.04 with Kernel 6.18.15-14-amd (x86_64), Insyde BIOS REM60070A, SMT Enabled, STAPM Disabled. - Intel Core Ultra X7 358H: MSI Prestige 16 Flip AI+ laptop C3MTG MS-2622, 32 GB LPDDR5X-8533 SK Hynix RAM, Kingston Technology KV3000 NVMe SSD 512 GB PCIe Gen4 512 GB storage, Linux Ubuntu 24.04 with Kernel 6.17.0-23-generic (x86_64), MSI BIOS E2622IMS.10E. Ryzen AI Max+ 395 (OPN 100-000002199) device is configured to reflect Ryzen AI Embedded X199 specifications. Intel data measured by AMD at 30W. 45W estimates were derived from AMD 30W measurements based on publicly available benchmark data. OEM published scores will vary based on system configuration and determinism mode used (default performance profile) and other factors. (REX-010)
4. Testing conducted by AMD as of May 2026 on an AMD Ryzen AI Max+ 395 (OPN 100-000002199) configured to reflect Ryzen AI Embedded X199 specifications (STAPM disabled; 45W SPL; sPPT 45W; fPPT 64W; LPDDR5X-8000), compared to an Intel Core Ultra X7 358H (30W; PL1 45W; PL2 64W; LPDDR5X-8533). Results reflect relative geometric mean performance across STREAM (GCC 15) on Ubuntu 24.04. System manufacturers may vary configurations, yielding different results. Results may vary based on configuration, settings, usage, and other factors. (REX-009)
5. Based on AMD internal testing as of July 2026, an AMD Ryzen AI Max+ 395 processor, configured to reflect Ryzen AI Embedded X199 specifications (AMD Maple CRB, 5.1 GHz CPU, 2.9 GHz GPU, with sustained operation at 45W TDP/sPPT/fPPT, and 64 GB soldered LPDDR5X-8000) was used to measure inference throughput. Results are compared against Intel Core Ultra X7 358H (MSI Prestige 16 Flip AI+ C3MTG MS-2622, 5.1 GHz CPU, 2.5 GHz GPU, PL1 = 45W, PL2 = 64W, MSI BIOS E2622IMS.10E, 32 GB LPDDR5X-8533, performance mode), and reflect the geometric mean of llama-bench (build: llama-b9453-vulkan) performance across gemma4 26B.A4B MXFP4 MoE, gemma4 26B.A4B Q4_K - Medium, llama3.1 8B Q4_K - Medium, qwen35 27B Q4_K - Medium, qwen35moe 35B.A3B MXFP4 MoE, qwen35moe 35B.A3B Q4_K - Medium models. All models fit in <24 GB RAM. System manufacturers may vary configurations, yielding different results. Results may vary based on configuration, settings, usage, and other factors. (REX-017)

DISCLAIMERS

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. CD-18u.

COPYRIGHT NOTICE

© 2026 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, Radeon, Ryzen, XDNA and combinations thereof are trademarks of Advanced Micro Devices, Inc. Core and Intel are trademarks of Intel Corporation. The terms HDMI, HDMI High-Definition Multimedia Interface, HDMI Trade dress and the HDMI Logos are trademarks or registered trademarks of HDMI Licensing Administrator, Inc. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. OpenGL and the oval logo are trademarks or registered trademarks of Hewlett Packard Enterprise in the United States and/or other countries worldwide. PCIe is a registered trademark of PCI-SIG Corporation. SPECrate is a trademark or registered trademark of Standard Performance Evaluation Corporation (SPEC). Ubuntu and the Ubuntu logo are registered trademarks of Canonical Ltd. Vulkan and the Vulkan logo are registered trademarks of the Khronos Group Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners. Certain AMD technologies may require third-party enablement or activation. Supported features may vary by operating system. Please confirm with the system manufacturer for specific features. No technology or product can be completely secure. PIDS146001