November 9, 2026

AMD EPYC Expands Its Leadership at the Foundation of AI

5th Gen AMD EPYC processors power Google Cloud TPU 8tx host infrastructure, helping enterprises protect sensitive AI models

News Highlights:

- *5th Gen AMD EPYC server CPUs power the host foundation for Google Cloud TPU 8tx, the first cloud Tensor Processing Unit (TPU) system with hardware-enforced confidential computing.*
- *AMD EPYC processors enable protection of sensitive AI models and data, combining host processing, security capabilities and x86 software compatibility.*
- *AMD EPYC processors are the host CPU of choice across more than 10 leading hyperscalers and frontier AI labs, spanning GPUs and custom AI accelerator ASICs.*

SANTA CLARA, Calif., Nov. 09, 2026 (GLOBE NEWSWIRE) -- AMD (NASDAQ: AMD) today announced that 5th Gen AMD EPYC™ processors are powering the host infrastructure for Google Cloud TPU 8tx, the first cloud Tensor Processing Unit (TPU) system to deliver hardware-enforced confidential computing. The expanded collaboration between the companies combines AMD EPYC™ CPU performance, AMD Infinity Guard security capabilities and x86 software compatibility with Google Cloud TPU AI acceleration to help enterprises deploy inference, fine-tuning and reinforcement learning workloads.[1] With Google Cloud's confidential computing architecture on TPU 8tx supported by 5th Gen AMD EPYC host processors, organizations can run their most sensitive AI workloads at scale in the cloud without compromising on data protection.

The deployment extends the role of AMD EPYC processors as a foundation for accelerated AI infrastructure. Across systems built with GPUs and custom accelerators, AMD EPYC CPUs are the choice for host processing – preparing data, coordinating workloads and managing communication with accelerators. With TPU 8tx, AMD and Google Cloud bring those capabilities together with hardware-rooted security to address a critical requirement for production AI: protecting data while it is in use.

"The most capable AI systems need a CPU foundation that can match the demands of the accelerators they host," said Dan McNamara, senior vice president and general manager, Compute and Enterprise AI, AMD. "Across generations of EPYC, we have built leadership capabilities in the performance, memory bandwidth and I/O that keep accelerated systems productive, backed by years of engineering collaboration with the world's leading cloud providers. Our work with Google Cloud brings that strength to TPU 8tx, helping customers scale AI with the performance and protection their most important workloads demand."

"When teams move AI agents into production, they need confidence that their data and intellectual property stay private, even from the underlying infrastructure," said Mark

Lohmeyer, vice president and general manager for AI and Computing Infrastructure, Google Cloud. “By pairing 5th Gen AMD EPYC processors with our 8th-generation TPU architecture inside the AI Hypercomputer, we are bringing hardware-enforced confidential computing to cloud TPUs for the first time. Customers get silicon-level data protection alongside the low-latency performance needed for modern reasoning models.”

Enabling Confidential AI with AMD EPYC Processors
In TPU 8tx systems, 5th Gen AMD EPYC processors manage workload orchestration, data preparation and high-throughput communication with TPU accelerators. The system utilizes the AMD Secure Encrypted Virtualization (SEV) capabilities in the AMD EPYC CPU, which together with the TPU, extends confidential computing protection beyond the CPU to connected accelerators, helping protect sensitive data as it moves among them.

Verifying an accelerator’s identity and security configuration before granting it direct access to protected memory helps customer prompts, model weights and intermediate data remain encrypted in CPU memory and in transit among the host and accelerators, with hardware isolation during accelerator execution. These protections help enterprises use sensitive information with confidence in reasoning and agentic AI applications.

AMD EPYC processors also underpin the broader Google Cloud confidential computing portfolio: AMD SEV technologies help protect workloads across Confidential VMs, Confidential GKE, Confidential Space, Confidential Managed Service for Apache Spark, and Confidential Dataflow. TPU 8tx extends that collaboration to confidential AI acceleration.

A Proven Host CPU Foundation Across AI Architectures
As hyperscalers, frontier labs and model providers build increasingly specialized AI infrastructure, CPU selection remains central to system performance and deployment. Accelerators depend on host processors to supply data, schedule work and connect AI workloads with the applications and services around them. AMD EPYC processors are the host CPU of choice for more than 10 leading hyperscalers and frontier AI labs, spanning GPUs and custom AI accelerator ASICs.

AMD EPYC processors support this role across systems using AMD Instinct™ GPUs, Nvidia GPUs and custom silicon. This breadth gives customers a proven CPU platform across accelerator choices, supported by the established x86 software ecosystem.

Building on Generations of AMD and Google Cloud Collaborations
AMD and Google Cloud have collaborated across multiple generations of EPYC processors, supporting general-purpose, compute-intensive, storage-intensive and accelerated workloads leading up to current Google Cloud Compute Engine C4D, N4D, H4D, Z4D and G4 virtual machines. The companies have also worked together for years on security technology at the corporate level, including deep reviews of AMD security processor firmware by Google security experts.

With TPU 8tx, this expanded collaboration helps enterprises run modern reasoning and agentic AI workloads in production without compromising on data protection.

Supporting Resources

- Learn more about AMD EPYC processors
- Explore AMD EPYC processors for GPU-accelerated systems
- Follow AMD on LinkedIn

- Follow AMD on X

About AMD
AMD (NASDAQ: AMD) drives innovation in high-performance and AI computing to solve the world's most important challenges. Today, AMD technology powers billions of experiences across cloud and AI infrastructure, embedded systems, AI PCs and gaming. With a broad portfolio of AI-optimized CPUs, GPUs, networking and software, AMD delivers full-stack AI solutions that provide the performance and scalability needed for a new era of intelligent computing. Learn more at www.amd.com.

AMD, the AMD Arrow logo, EPYC, Instinct 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] GD-183A: AMD Infinity Guard features vary by EPYC™ Processor generations and/or series. Infinity Guard security features must be enabled by server OEMs and/or Cloud Service Providers to operate. Check with your OEM or provider to confirm support of these features. Learn more about Infinity Guard at https://www.amd.com/en/technologies/infinity-guard

Contact:
Aaron Grabein
AMD Communications
aaron.grabein@amd.com

Liz Stine
AMD Investor Relations
liz.stine@amd.com

Source: Advanced Micro Devices, Inc.