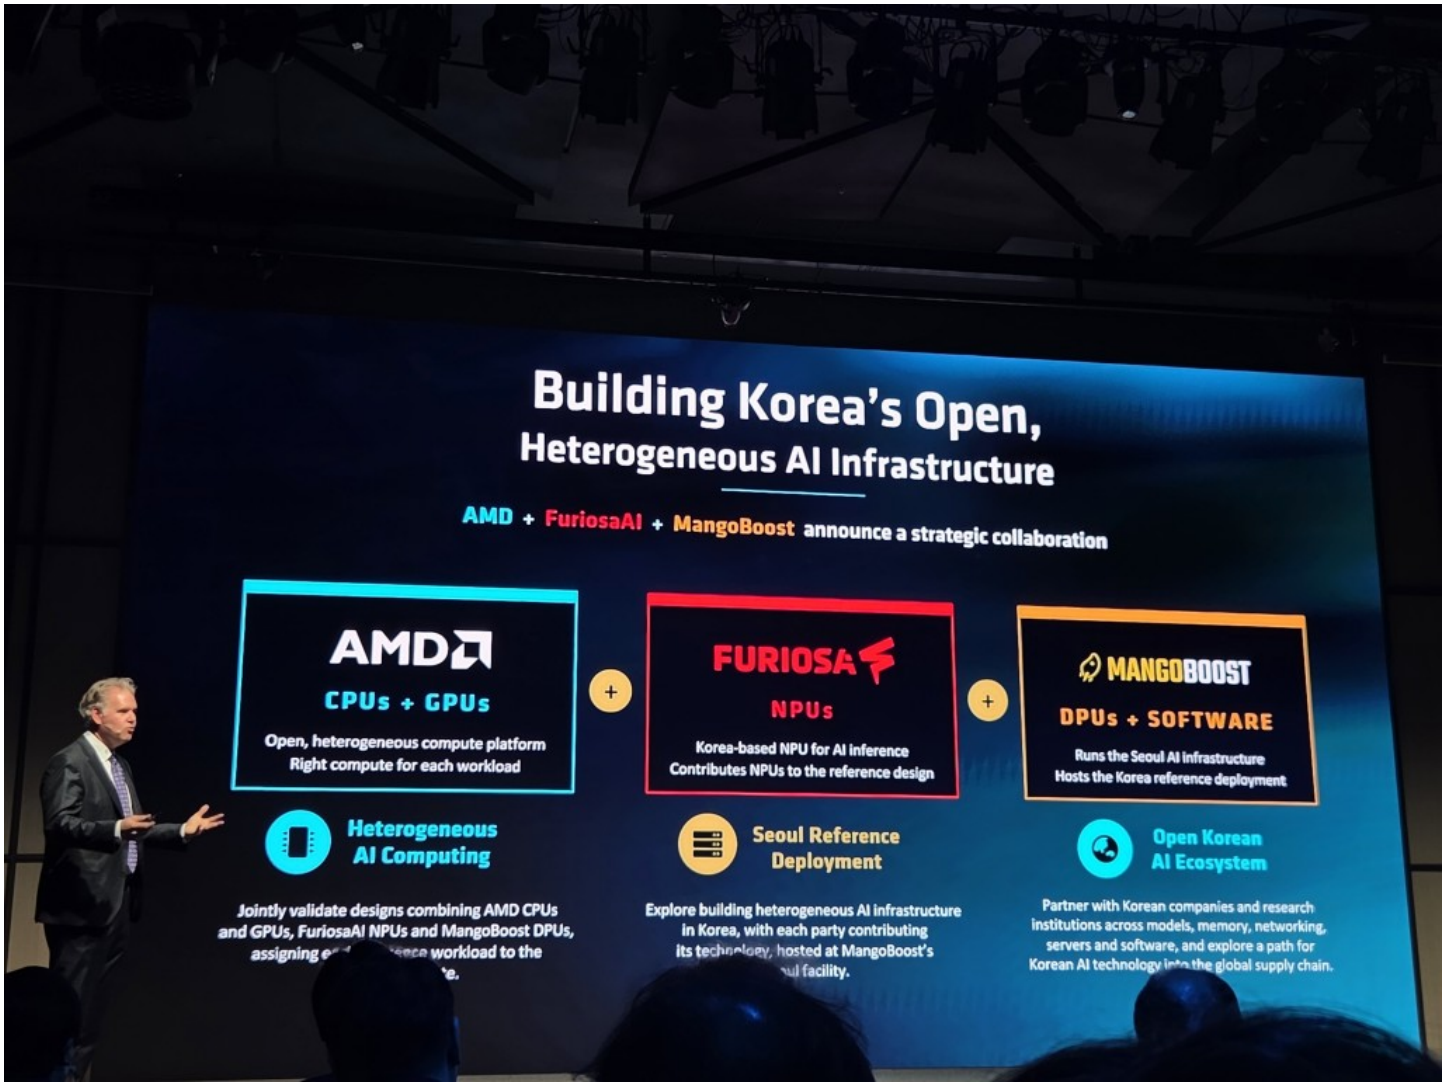
MangoBoost Signs MoU with AMD and FuriosaAI to Explore Next-Generation AI Infrastructure

Three companies announce MOU at AMD Advancing AI Seoul 2026, strengthening full-stack AI infrastructure capabilities centered on the Mango AI Factory



Seoul, South Korea Nov 3, 2026 (Issuewire.com) -

- MangoBoost, AMD, and FuriosaAI to explore open, heterogeneous AI infrastructure integrating CPUs, GPUs, NPUs, and DPUs
- MangoBoost advances development of its Korea-based Mango AI Factory powered by AMD Instinct GPUs, leveraging full-stack optimization to enhance performance and operational efficiency
- MangoBoost CEO Jangwoo Kim highlights the industrial value of AI infrastructure and showcases Mango Inference™ at the event

MangoBoost, a full-stack AI infrastructure company led by CEO Jangwoo Kim, announced that it has

signed a Memorandum of Understanding (MOU) with AMD and FuriosaAI at the AMD Advancing AI Seoul 2026 & AI DevDay Seoul 2026 to collaborate on the development of next-generation open, heterogeneous AI infrastructure. The MoU marks a step toward strengthening MangoBoost's AI infrastructure business capabilities, with a focus on its Korea-based Mango AI Factory.

The three companies intend to jointly explore and validate infrastructure designs combining AMD CPUs and GPUs, FuriosaAI NPUs, and MangoBoost's DPU technology and system software optimization capabilities to enable high-performance, efficient infrastructure capable of supporting diverse AI workloads. They will also explore a reference deployment using MangoBoost's AI infrastructure in Seoul, with MangoBoost targeting the first hybrid reference infrastructure to become operational in 2027.

Through this collaboration, MangoBoost aims to extend its proprietary BoostX™ DPU-based network acceleration technology and system software optimization capabilities across diverse AI accelerator environments. The company plans to further strengthen its technological competitiveness by improving system performance and resource utilization efficiency in heterogeneous computing environments.

In parallel, MangoBoost is advancing the development of its Korea-based Mango AI Factory, targeting deployment in the first half of 2027. Powered by AMD Instinct™ GPUs, the Mango AI Factory is designed as a full-stack AI infrastructure platform integrating high-performance GPU computing resources with MangoBoost's BoostX™ DPU, LLMBoost™ AI inference optimization software, Alphonso™ GPU servers, and Kesar™ storage solutions.

As large-scale AI models and agentic AI applications continue to expand, AI service performance increasingly depends not only on GPU compute capabilities but also on the efficiency of data movement, networking, and memory processing. MangoBoost's strategy is to optimize hardware and system software as an integrated stack, reducing infrastructure bottlenecks, improving GPU utilization and AI inference efficiency, and lowering total cost of ownership (TCO).

Beyond supplying GPU computing resources, MangoBoost aims to deliver an end-to-end business model spanning infrastructure deployment, system optimization, and AI inference services. This integrated approach is designed to address enterprise requirements for performance and cost efficiency while supporting data residency and control through GPU clusters located in Korea.

Meanwhile, MangoBoost CEO Jangwoo Kim participated in a fireside chat titled “Building Korea's AI Nation: From Infrastructure to Impact,” where industry and government leaders discussed strategies for strengthening Korea's AI infrastructure competitiveness and accelerating real-world industrial adoption.

Moderated by Jeongnghee Jang, Head of Commercial Sales at AMD Korea, the session brought together Jangwoo Kim, CEO of MangoBoost; Kangwon Cho, CEO of Moreh; and Dogyu Lee, Director General for ICT Policy at the Ministry of Science and ICT. The discussion focused on strengthening Korea's AI computing capabilities, expanding developer access, and addressing key challenges in translating AI infrastructure investments into practical industrial applications.

“AI infrastructure competitiveness is defined not simply by the scale of computing resources, but by how efficiently those resources are utilized and translated into real-world service value,” said Jangwoo Kim, CEO of MangoBoost. “By combining AMD Instinct GPUs with our proprietary full-stack optimization technologies, we aim to deliver high-performance, efficient AI infrastructure through the Mango AI Factory while contributing to the growth of an open, heterogeneous AI infrastructure ecosystem.”

At the event, MangoBoost also demonstrated Mango Inference™, its serverless AI inference platform powered by AMD Instinct GPUs and the company's proprietary LLMBoost™ AI inference optimization software.

Mango Inference™ enables developers and enterprises to access a wide range of open-weight AI models through APIs without the need to build or manage GPU infrastructure. Designed to support diverse AI workloads, including coding agents, long-context processing, and advanced reasoning, the platform focuses on delivering high throughput, low latency, and a streamlined developer experience.

Looking ahead, MangoBoost plans to integrate the Mango AI Factory with Mango Inference™ to further strengthen its full-stack AI business capabilities, spanning infrastructure deployment, AI model execution, and service operations. Building on this foundation, the company intends to expand AI infrastructure and inference service opportunities across global markets, including the Americas, Japan, and the Middle East, while deepening its technology partnerships.

About MangoBoost

MangoBoost is a full-stack AI infrastructure company focused on maximizing AI infrastructure performance and efficiency. Its portfolio includes BoostX™ DPU-based network acceleration technology, LLMBoost™ AI inference optimization software, Alphonso™ GPU servers, Kesar™ storage solutions, and Mango Inference™, a serverless AI inference platform.

By optimizing hardware, system software, and AI services as an integrated stack, MangoBoost enables high-performance, efficient AI infrastructure designed to meet the evolving demands of modern AI workloads.

Media Contact

contact@mangoboost.io





Media Contact

MangoBoost

*********@mangoboost.io

https://www.mangoboost.io/

Source : MangoBoost

See on IssueWire