Globe Newswire (2026-01-19)

Chinese RISC-V Chipmaker SpacemiT Launches K3 AI CPU, Highlighting the Rise of Open-Source Hardware in Intelligent Computing

Hangzhou, China, Jan. 29, 2026 (GLOBE NEWSWIRE) -- As the global technology industry accelerates its shift toward open architectures and on-device artificial intelligence, Chinese RISC-V chip company SpacemiT announced the launch of its new K3 AI CPU on January 29. The company aims to combine the open RISC-V instruction set with general-purpose and AI computing capabilities, offering a more flexible, power-efficient and cost-effective platform for intelligent terminals and edge computing.

For decades, the processor market has been dominated by x86 and Arm architectures. However, as AI workloads gradually move from the cloud to end devices, there is growing interest in more open and customizable computing platforms. RISC-V, as a fully open-source instruction set architecture, is increasingly seen as a key foundation for the next generation of open computing ecosystems.

Founded in 2021 and headquartered in Hangzhou, SpacemiT is one of the few Chinese semiconductor companies committed to a “pure RISC-V” strategy. The company focuses on integrating high-performance general

computing and AI acceleration into a single chip, which it describes as an “AI CPU” approach. This design philosophy targets intelligent hardware scenarios that require high computing density, low power consumption and strong system integration.

The K3 chip is the result of more than 1,200 days of development. According to the company, it is among the first mass-production-ready RISC-V AI CPUs compliant with the RVA23 specification. It also supports 1024-bit RISC-V Vector extensions (RVV) and native FP8 precision for AI inference.
In terms of hardware configuration, K3 integrates eight high-performance X100 RISC-V CPU cores with a maximum frequency of 2.4GHz. SpacemiT said its single-core performance is broadly comparable to Arm’s Cortex-A76.

The chip delivers up to 60 TOPS of AI compute and supports up to 32GB of LPDDR5 memory. SpacemiT noted that K3 is not designed to compete directly with high-end server CPUs or GPUs, but instead to enable local execution of medium-scale AI models and multimodal applications. The company said K3 can support models in the 30- to 80-billion-parameter range on a single chip, with typical system power consumption between 15 and 25 watts.

On the software side, K3 adopts a co-design approach between hardware and software. It supports mainstream AI frameworks and compilers such as Triton and TileLang, and is compatible with major open-source AI

ecosystems and Linux distributions. The company said this is intended to reduce the development barriers for deploying AI models on RISC-V platforms, bringing the experience closer to that of x86 and Arm systems.

“We believe the long-term direction of computing architectures is a transition from closed to open systems,” said Chen Zhijian, founder and CEO of SpacemiT, at the launch event. “x86 is highly closed, Arm is semi-open, while RISC-V is fully open. In the long run, open instruction sets are more likely to become the foundation of global computing.”

Chen added that RISC-V carries particular significance for China’s semiconductor industry. “In the past, Chinese computing chips were largely limited to domestic markets. Open architectures create a new path for Chinese chips to integrate more naturally into the global technology ecosystem.”

K3 emphasizes a fusion of general computing, AI computing and data-coherent interconnects. Chen described this direction as the evolution toward a “next-generation AI CPU,” where traditional CPUs become increasingly intelligent.

“In the AI era, CPUs can no longer be just control processors,” he said. “They must also provide native AI computing capabilities. This is similar to the shift from feature phones to smartphones — a fundamental change in the form of computing.”

SpacemiT also disclosed that its previous-generation K1 chip has achieved shipments of more than 150,000 units and has been deployed in industrial control systems, robotics, edge computing platforms and open-source intelligent hardware. This commercial experience provides a foundation for the rollout of K3. The company said K3 has already received orders, with initial deliveries planned from the end of April 2026.

The company continues to promote a “full-stack RISC-V” strategy, covering CPU IP, chip design, operating systems, compilers, AI software stacks and developer platforms. K3 supports multiple operating systems including

Ubuntu, OpenHarmony and OpenKylin. SpacemiT is also launching supporting products such as PICO-ITX single-board computers, robot core boards and array server platforms, while opening its hardware reference designs to developers and system integrators.

Hangzhou has recently emerged as a major hub for AI, semiconductor and open-source technology innovation in China. Industry observers note that a new generation of Chinese technology companies is taking shape in the region, strengthening China's presence in advanced "hard-tech" sectors. SpacemiT is regarded as one of the representative players in this ecosystem.

At the same time, the industry generally acknowledges that RISC-V still lags behind x86 and Arm in high-end computing, software ecosystem maturity and overall industrial scale. SpacemiT also recognizes that its current products are better suited to intelligent terminals and edge AI rather than competing directly with top-tier server processors or GPUs.

"Our goal is not to confront global giants head-on," Chen said. "Instead, we aim to establish differentiated advantages in the mid-range computing segment, using lower power consumption, higher integration and better cost efficiency to make AI computing accessible to more devices."

As artificial intelligence continues to move from centralized cloud platforms toward local deployment, power efficiency, system integration and open ecosystems are becoming key competitive factors. The launch of K3 represents a concrete step by China's RISC-V community in combining open-source architectures with AI computing, and reflects China's broader effort to explore new technological paths in the next generation of global computing.