AMD

WHY IS EVERYONE TALKING ABOUT CPUs AND AGENTIC AI?

CPU demand is exploding thanks largely to agentic AI — here are five reasons why

After years of massive GPU buildouts, tech leaders are suddenly talking about CPUs and how AI needs a lot more of them. Why the new focus on CPUs?

Agentic AI is making CPUs more relevant because it multiplies workloads that CPUs do best. As agentic AI adoption explodes, many hyperscalers and businesses are turning to AMD EPYC™ Server CPUs for the high core and thread counts, high performance, and large memory capacity and bandwidth that agentic AI workloads need.

Agentic AI is just getting started. Read on to see how it will drive unprecedented demand for CPUs.

1 ENTERPRISES ARE RAPIDLY ADOPTING AGENTIC AI

Agentic AI can assist virtually any business role that involves a computer. Employees use AI agents to write code, build websites, and more. Every employee may make dozens of requests from dozens of agents an hour. That means a company with 1,000 employees could conceivably launch hundreds of thousands of agentic requests a day.

- 74% of companies plan to deploy agentic AI in the next two years.[1]

2 AGENTIC AI WORKFLOWS RUN ON CPUs

Agentic AI workflows — orchestrating, calling tools, requesting inference, executing instructions — run on the CPU side. Each agentic AI request can create dozens of CPU-executed tasks. As agentic AI adoption grows, more CPUs must come online to meet the demand.

- CPU-based tool processing and orchestration account for 50–90% of the total end-to-end runtime for agentic AI pipelines.[2]

3 ONE AGENTIC AI REQUEST CAN SPAWN A DELUGE OF AI AGENTS

When a person asks an AI agent to do something complex like building a travel itinerary, an orchestration agent reviews the request, breaks it into tasks, and assigns those tasks to multiple subagents. Those subagents then spin up more subagents to search, run APIs, and assemble the finished itinerary. Multiply that by all the agentic requests made per second around the world, and you have a recipe for massive numbers of agents and immense growth in CPU workloads.

- A major SaaS provider reports it has 3,258 primary AI agents working alongside its 1,300 employees[3], a ratio of 2.5 to 1.

AMD

4 SEARCH ENGINES, DATABASES, CRM, ERP — EVERYDAY SYSTEMS MUST SCALE UP TO MEET AGENTIC DEMAND

Agentic AI uses the same tools, applications, and systems that people use to find information and complete tasks – with one major difference. A person can only use a few tools at a time. One agentic project can run multiple agents in parallel, creating cascading API calls to business systems and software – that all typically run on CPUs. To meet demand and maintain service levels, providers and businesses will have to add CPUs to their fleets.

- AI agents now create 80% of databases and 97% of test and development environments on the Databricks platform.[4]

5 AGENTIC AI DRIVES MORE INFERENCE

Each agentic request can trigger multiple calls to LLMs for inference. Orchestration agents call models to plan and route work. Subagents call models to generate code, content, and queries. This explosion in inference calls increases demand for GPU cycles, which drives more demand for CPUs. Because every GPU depends on CPUs to function.

- Inference workloads have grown from roughly one-third of all compute in 2023 to one-half in 2025. In 2026, inference is expected to consume roughly two-thirds of all compute.[5]

WHILE EVERYONE ELSE IS TALKING ABOUT CPUs, *AMD IS DELIVERING THEM*

Agentic AI is already driving massive demand for CPUs. Based on that surge, AMD now projects the total addressable market (TAM) for CPUs to roughly quadruple by 2030.[6]

AMD is helping to drive the next phase of AI infrastructure with our 5th Generation AMD EPYC™ Server CPUs. Even more workload-specific performance for agentic AI will be possible with our 6th Generation AMD EPYC™ Server CPUs, slated to be available in the second half of 2026.

Learn more about agentic AI

1. Deloitte, The State of AI in the Enterprise, Copyright © 2026 Deloitte Development LLC.
2. Ritik Raj, Hong Wang, Tushar Krishna, A CPU-Centric Perspective on Agentic AI, November 2025.
3. Zeb Evans, Founder and CEO of ClickUp, How Many Agents Work at Your Company? (LinkedIn post).
4. Databricks, 2026 State of AI Agents.
5. Deloitte, Why AI's next phase will likely demand more computational power, not less, Copyright © 2025 Deloitte Development LLC.
6. AMD, AMD Financial Analysts Day, November 11, 2025 and Q1 2026 AMD Earnings Call.

© 2026 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and other countries. Other product names used herein are for identification purposes only and may be trademarks of their respective owners.

AMD
together we advance_