


Qualcomm AI Server Motherboard Model



Overview

From Qualcomm's vantage point, the AI200 and AI250 deliver rack-scale performance and superior memory capacity at an industry-leading total cost of ownership (TCO), with optimized performance for large language models (LLMs) and large multimodal models (LMMs). The cards interface to the host via PCIe and to the BMC via SMBus/I2C. The Cloud AI 100 Ultra SKU has 4. Qualcomm is moving from mobile NPUs into rack-scale AI infrastructure, positioning its AI200 (2026) and AI250 (2027) to challenge Nvidia/AMD on the economics of large-scale inference. Qualcomm is entering this evolving market with its Dragonfly platform, aiming for \$15 billion in revenue by 2029. The stock soared 11% following the news.



Article Content

Qualcomm Unveils AI200 and AI250: Rack-Scale

The most disruptive part is Qualcomm AI250, debuting with an architecture based on near-memory computing. Its declared goal: a generational leap in efficiency

Qualcomm Unveils AI200 and AI250—Redefining Rack-Scale Data

With seamless compatibility for leading AI frameworks and one-click model deployment, Qualcomm AI200 and AI250 are designed for frictionless adoption and rapid innovation.”

Qualcomm Launches On-Prem AI Appliance Solution

The Qualcomm AI On-Prem Appliance Solution is powered by the Qualcomm® Cloud AI roadmap of accelerators. It combines the accessibility and

Qualcomm Hexagon NPU | Snapdragon NPU Details

The Qualcomm® Hexagon™ NPU, purpose-built for AI, powers advanced gen AI models and cutting-edge gen AI experiences while maintaining best-in-class

Cerebras Enables Even Larger Generative AI Models

To offer customers a solution to transition from training complex traditional and generative AI models to efficiently inference processing those

Architecture

The seventh-generation AI Core, leveraging over a decade of Qualcomm AI research, is shown in Figure below. The AI Core implements the architecture

Qualcomm® AI Hub Models (Compute)

The platform for on-device AI, with optimized open source and licensed models, or bring your own. Validate performance on real Qualcomm devices.

A Look At Qualcomm's Data Center Inference Accelerator

Cloud AI 100 Qualcomm's first series of inference accelerators is called the Cloud AI 100. “We are currently actively collaborating with several industry leaders for deploying this solution,” said

Qualcomm's AI Data Center Bet: Inside The Dragonfly Strategy

Qualcomm structured its data center strategy around four product lines that layer on top of one another, mirroring the approach it used in automotive, where Qualcomm entered through...

Best Motherboards for AI and Machine Learning Builds

Explore the best motherboards for AI and machine learning builds in 2025. Compare PCIe lanes, GPU support, RAM capacity, and advanced features.

Deep Dive Into Qualcomm's Centriq Arm Server Ecosystem

Reference Motherboard And Platform Enable System Ecosystem Qualcomm's Centriq reference motherboard design is a half-width board designed to fit into a wide range of 1U bent-metal

Qualcomm announces AI chips to compete with AMD

AI labs need that computing power to run the most advanced models. Qualcomm's data center chips are based on the AI parts in Qualcomm's

And so it begins: Gigabyte quietly launched an AI

Pro And so it begins: Gigabyte quietly launched an AI motherboard with built in surge protection and support for four dual-slot GPUs — so could AI

Qualcomm Cloud AI100 Ultra: 4 Times The

Now Qualcomm Cloud AI100 can handle 100B models on the data center inference processor to help improve LLM inference affordability at

Qualcomm Launches AI250 & AI200 With Huge

From Qualcomm's vantage point, the AI200 and AI250 deliver rack-scale performance and superior memory capacity at an industry-leading total

Snapdragon Processors = Extraordinary Experiences

Snapdragon processors give your smartphone the ability to create extraordinary experiences in camera, gaming, and audio. Learn how we do it and

Advantech Qualcomm-Based Solutions

Advantech is redefining the edge with an innovative portfolio of Qualcomm-based solutions, designed to meet the evolving demands of diverse industries. Our offerings include scalable pre-configured AI

Qualcomm's AI Data Center Bet: Inside The Dragonfly Strategy

Qualcomm's data center entry analyzed: how the Modular acquisition, HBC memory architecture, and Arm-based C1000 CPU challenge NVIDIA in AI inference.

Qualcomm enters data center market with new AI

Qualcomm is making its debut in the data center hardware market with two new AI accelerator chips, the AI200, set for release in 2026, and the

Qualcomm AI200/AI250 AI Chips for Data Center

Evaluate Qualcomm's AI200/AI250 against workload fit, software portability, facility readiness, and multi-vendor risk posture. Map your 24–36

Qualcomm Enters AI Server CPU Market: Challenges X86 with Arm ...

Select Language: Qualcomm Has Entered the AI Server CPU Market, Challenging x86 Architecture with Arm-Based Solutions In a significant move within the tech industry, Qualcomm has

Qualcomm AI Hub

The platform for on-device AI, with optimized open source and licensed models, or bring your own. Validate performance on real Qualcomm devices.

Qualcomm® Cloud AI 100 Ultra

Qualcomm Cloud AI 100: Leading in all inference performance metrics with industry-leading advancements in performance density and performance-per-watt capabilities.

Qualcomm announces AI chips to compete with AMD

Qualcomm said that both the AI200, which will go on sale in 2026, and the AI250, planned for 2027, can come in a system that fills up a full, liquid

Qualcomm Unveils Dragonfly Data Center Roadmap: C1000 CPU,

Dragonfly C1000 CPU and Data Center Platform Qualcomm introduced the Dragonfly C1000, a purpose-built data center CPU designed around its custom Oryon cores, targeting agentic

Cloud AI 100 Ultra Processor | Generative AI Leader

The Qualcomm Cloud AI 100 Ultra is groundbreaking: with the ability to deploy a staggering 100 billion parameter model on a single-width, energy

AI inference at scale with AI250 rack-scale Platform | Dragonfly

Qualcomm Dragonfly AI250 rack-scale AI inference platform delivering 768GB memory, 133TB/s bandwidth for low-latency, cost-efficient real-time AI.

Edge AI Computing | On-Premise AI Acceleration

Transforming business with on-prem AI The Qualcomm Dragonwing™ AI On-Prem Appliance Solution enables businesses to run custom and off-the-shelf AI

Contact Us

For more information, pricing, or custom solutions, please contact us:

Website: <https://www.sphplantsales.co.za>

Email: info@sphplantsales.co.za

Phone: +27 64 827 3915

Address: Unit 9, Highveld Technopark, 43 Atlas Road, Johannesburg, 2196,
South Africa

This document is for informational purposes only. Specifications subject to
change without notice.

