CORSAIR AI WORKSTATION 300
DESKTOP PC

COMPRESS ED POWER

Designed for efficiency and adaptive performance, it's ideal for AI developers, creators, and professionals seeking a powerful streamlined AI workstation solution. Powered by AMD Ryzen™ AI Max 300 Series—ready for local LLMs, creative tasks, and AI development.

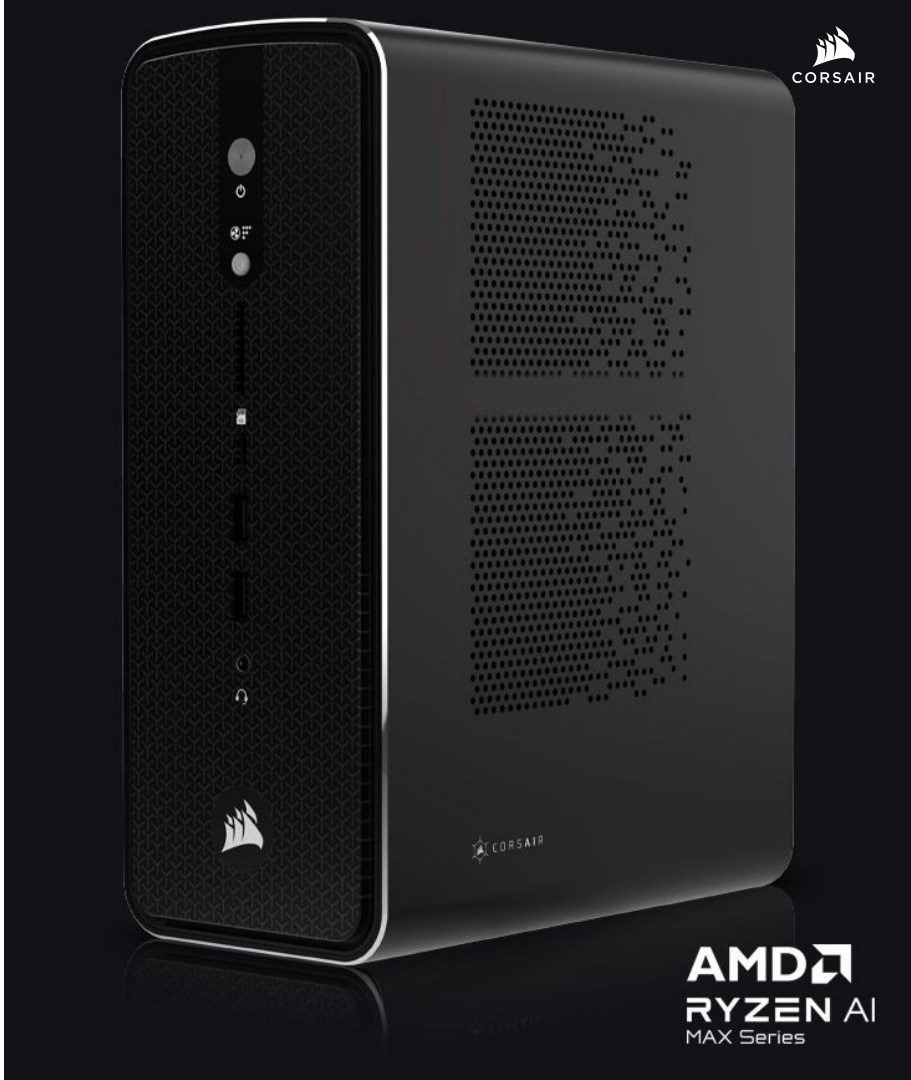


Power Button
SD 4.0 Card Reader
2x USB 3.2 Gen 2 Type-A Ports
Performance Level Selector
1x USB 4.0 Type-C Port
Headphone / Mic Combo Jack
2x USB 2.0 Type-A Ports
DisplayPort 2.1
1x USB 3.2 Gen 2 Type-A Port
HDMI 2.1 Port
1x USB 4.0 Type-C Port
Ethernet Port
Headphone / Mic Combo Jack

COMPACT DESIGN

Experience AI performance out of the box, designed for quiet efficiency and modular flexibility.

AI-READY PERFORMANCE

Accelerate your AI workflows with a high performance APU & dynamic VRAM for local LLMs, creative tasks, and development.

ENGINEERED FOR SECURITY

Layers of built-in security technology for chip-to-cloud protection against sophisticated attacks.

CHOOSE YOUR WORKSTATION

Starting at: $1,599 or as low as $96.09/month	Starting at: $1,999 or as low as $120.12/month	Starting at: $2,299 or as low as $138.15/month
		
AMD Ryzen™ AI Max 385 (8C/16T)	AMD Ryzen™ AI Max+ 395 (16C/32T)	AMD Ryzen™ AI Max+ 395 (16C/32T)
64GB LPDDR5X-8000MT/s	128GB LPDDR5X-8000MT/s	128GB LPDDR5X-8000MT/s
1TB PCIe NVMe	1TB PCIe NVMe	4TB (2x 2TB) PCIe NVMe
AMD Radeon 8050S up to 48GBs VRAM	AMD Radeon 8060S up to 96GBs VRAM	AMD Radeon 8060S up to 96GBs VRAM
2-Year Warranty	2-Year Warranty	2-Year Warranty

Corsair Ai 300 Mini PC based on AMD’s Strix Halo platform for Ai Developers delivers the Same Power for Half the Price of NVIDIA’ DGX Spark ‘So-Called’ Supercomputer —

VS.

128GB LPDDR5x
Unified RAM

$1,999

$3,999 +

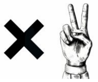

Build a 2-Node Cluster for LLMs for the same price of 1 DGX Spark

NVIDIA's DGX Spark was a monumental release from the firm, but at the same time, Corsair collaborate with AMD to create a platform delivering high performance on-device AI with Strix Halo APU to perform better than NVIDIA's GB10 chip in various AI workloads.

NVIDIA's DGX Spark Is a Solid Option For Throughput, But Strix Halo Is the Go-To Platform For the Best Price-to-Perf Figures

Despite featuring an impressive performance, consumers have voiced their opposition to the price tag of the DGX Spark, claiming that the $4,000 price tag makes the DGX Spark less attractive. However, Corsair launched one of AMD's most reputable AI platform offers an impressive alternative to the DGX Spark for almost half the price.

AMD Strix Halo Vs. NVIDIA DGX SPARK:

Local AI Performance Benchmark

Test Model	Metric	AMD Strix Halo	NVIDIA GB10	Winner
Llama 3.3 70B	Generation Speed (tok/sec)	4.9	4.67	AMD
	First Token Response Time (s)	0.86	0.53	NVIDIA
Qwen3 Coder	Generation Speed (tok/sec)	35.13	38.03	NVIDIA
	First Token Response Time (s)	0.13	0.42	AMD
GPT-OSS 20B	Generation Speed (tok/sec)	64.69	60.33	AMD
	First Token Response Time (s)	0.19	0.44	AMD
Qwen3 0.6B Model	Generation Speed (tok/sec)	163.78	174.29	NVIDIA
	First Token Response Time (s)	0.02	0.03	AMD

Performance Insights

Across the tests, Corsair **Ai 300** demonstrates remarkable efficiency in handling local AI workloads.

- For **medium to large models (20B–70B)**, the Ai 300 consistently delivers faster generation speeds, particularly excelling in the **GPT-OSS 20B** and **Llama 3.3 70B** benchmarks.
- In **first-token latency**, Ai 300 achieves significantly lower response times in multiple models, showcasing AMD's optimized memory bandwidth and local AI acceleration architecture.

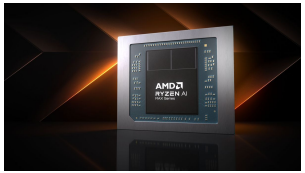

🚀 Why Corsair Ai 300 Excels in Local AI

1. **AMD Ryzen AI Engine**
   Harnesses dedicated AI acceleration for multi-threaded, low-latency inference performance.
2. **Open Ecosystem Compatibility**
   Fully supports major open-source AI models (Llama, Qwen, Mistral) and frameworks (Ollama, LM Studio, vLLM).
3. **Superior Performance-to-Power Ratio**
   Delivers workstation-level AI capability within a compact Mini PC form factor, with exceptional energy efficiency.
4. **Local AI Workflow Ready**
   Deploy and run models directly on-device — no cloud dependency, enhanced data privacy, and instant responsiveness.