



“ At Arctur, the largest privately owned HPC and AI infrastructure provider in Central Europe, we are already using the MACx AI accelerators to power LLM inference in production — with **outstanding price-performance, great efficiency, and supply independence.**

Tomislav Subić, Head of AI, Arctur

Test before you buy - production environment available

1 node (4 GPUs), 24h — €100

1 node (8 GPUs), 24h — €200

Maximum testing capacity: 10 nodes

Testing cost is fully credited

toward purchase

Toraks d.o.o.

sales@toraks.si

www.toraks.si

MACx GPU Accelerator

Model: MACx-01 PCIe AI Accelerator

Scale AI Without Vendor Lock-In

Enterprise-grade AI inference acceleration built for production LLM workloads and **scalable data center infrastructure.**



Why MACx?

- 64GB HBM2e high-bandwidth memory
- Up to 340 TFLOPS AI performance
- PCIe 5.0 x16 interface
- Optimized for production AI inference
- Native PyTorch and TensorFlow support

Deployment & Pricing

List price: €5,490 per unit

Promotional pricing:

- Minimum purchase: 4 units - 4,990 per unit
- 16 units - 4,790 per unit
- 64 units - 4,650 per unit
- For larger volumes - please contact us
- Warranty: 2 years standard

Technical Specifications

Performance

FP64: 21 TFLOPS

FP32: 42 TFLOPS

TF32: 170 TFLOPS

BF16: 340 TFLOPS

FP16: 340 TFLOPS

INT8: 680 TOPS

Video Decode

90 streams H.264 / VP9

80 streams H.265

110 streams AVS2

Memory

Capacity: 64GB HBM2e

Bandwidth: 1.8 TB/s

High-bandwidth memory optimized for AI inference workloads

Interface

PCIe 5.0 x16

Form Factor

Full-height, full-length PCIe card

Power

400W air-cooled