

OpenAI Jalapeño

Jalapeño

Custom

2026

OVERVIEW

— FP32 TFLOPS	216 GB Memory HBM4	15.4 TB/s Memory Bandwidth	700 W TDP
------------------	-----------------------	-------------------------------	--------------

PERFORMANCE METRICS

Precision	Peak (dense)	Bits	Type
FP64	—	—	Standard
FP32	—	—	Standard
TF32	—	—	Standard
BF16	—	16	Standard
FP16	—	—	Standard
FP8	3400.0	8	Standard
FP6	—	—	Standard
FP4	13400.0	4	Standard
INT8	—	—	Standard

MEMORY SPECIFICATIONS

Capacity	216 GB
Type	HBM4
Bandwidth	15.4 TB/s
Interface Width	—

POWER & EFFICIENCY

TDP	700 W
Max power (Flopper estimate)	~805 W est.
FP32 dense TFLOPS per W of TDP	—
FP16 dense TFLOPS per W of TDP	—
Flopper estimate: 700 W TDP x 1.15. The vendor publishes no maximum board power for this part. Per W of TDP figures divide dense peak by accelerator TDP. Board power only; excludes host CPUs, networking, cooling and facility overhead.	

HARDWARE DETAILS

Vendor: OpenAI	Architecture: Jalapeño	Process Node: —
Form Factor: Custom	Launch Year: 2026	

INTERCONNECT

GPU-to-GPU: Ethernet scale-up
(Broadcom Tomahawk 6)

Bandwidth: 600 GB/s
direction and scope not stated by vendor

CHIP DESIGN

Chiplets: 2 (1 compute die + 1 IO
chiplet)
