


# Tensix Processor PCIe Cards

Open • Scalable • Performant



Tenstorrent's Grayskull™ and Wormhole™ Tensix Processors are open, flexible, and scalable, offering superior performance for cost compared to traditional GPUs. The processor is comprised of Tensix Cores, each of which includes local scratch pad SRAM, five “baby RISC-V” cores, and Matrix and Vector Engines with broad precision format support. Wormhole™ processors feature a Network-on-Chip and use Ethernet protocols to network into a multichip mesh for workstations and servers (such as Galaxy). All Tenstorrent Tensix Processors are supported by two open source SDKs for either high-level (TT-Buda™) or low-level (TT-Metalium™) development.

## Comparison Chart

Card	e75	e150	n150s	n300s
Part Number	TC-01001	TC-01002	TC-02001	TC-02003
ASIC	Grayskull™	Grayskull™	Wormhole™	2x Wormhole™
Tensix Cores	96	120	72	128 (64 per ASIC)
AI Clock	1 GHz	1.2 GHz	1 GHz	1 GHz
SRAM	96MB	120MB	108MB	192MB (96MB per ASIC)
Memory	8GB LPDDR4	8GB LPDDR4	12GB GDDR6	24GB GDDR6
Memory Speed	3.2 GT/sec	3.7 GT/sec	12 GT/sec	12 GT/sec
Memory Bandwidth	102.4 GB/sec	118.4 GB/sec	288 GB/sec	576 GB/sec
TFLOPs (FP8)	221	332	262	466
TFLOPs (FP16)	55	83	74	131
TFLOPs (BFP8)	55	83	148	262
Interface	PCIe 4.0 x 16	PCIe 4.0 x 16	PCIe 4.0 x 16	PCIe 4.0 x 16
Internal Interconnect	-	-	-	400GbE
Total Board Power	75W	200W	160W	300W
Cooling	Active (Blower)	Passive*	Passive*	Passive*
Form Factor	Single Slot, HHHL	Dual Slot, FHFL	Dual Slot, FHFL	Dual Slot, FHFL

\*Active Cooling Kit (TX-01001) included.

## Flexible Precision Support

Tenstorrent's Tensix Cores support a broad range of data types, including highly efficient block floating point (BFP) precision. BFP offers most of the precision of conventional floating point formats while requiring just half the bandwidth and storage.

	Grayskull™ e75/e150	Wormhole™ n150s/n300s
Floating point	FP8, FP16, BF16	FP8, FP16, BF16, FP32*
Block floating point	BFP2, BFP4, BFP8	BFP2, BFP4, BFP8
Integer	-	INT8, INT32*
Unsigned integer	-	UINT8
TensorFloat	-	TF32
Vector	VTF19	VFP32, VTF19

\*Output only.

## Open Source SDKs

Tenstorrent's Tensix Processors are supported by two open source SDKs. TT-Buda™ uses a top-down approach, compiling code from common ML frameworks like PyTorch or TensorFlow and abstracting the underlying hardware. TT-Metalium™ has a bottom-up approach that provides access directly to the metal, enabling use of Python and C++ for a variety of workloads.

