TENSORDYNE

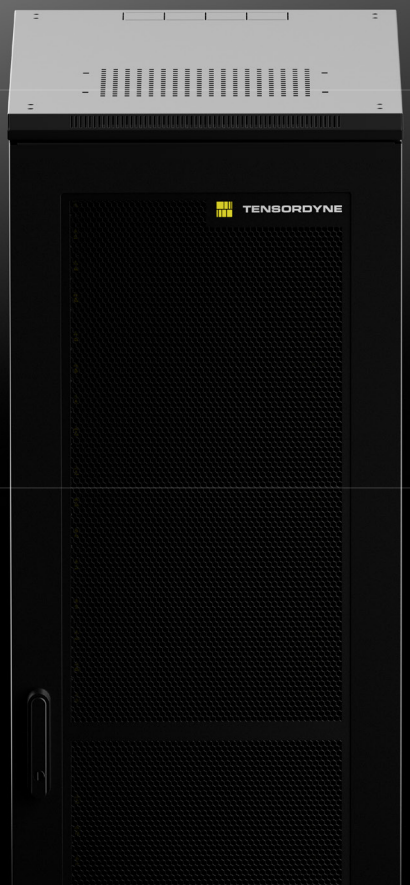

Napier Generation
AI Inference
Compute System

TDN AIP
TDN ACT
TDN LINK
TDN SDK
TDN RACK

Tensordyne Napier (TDN) Datasheet

JUNE 2026 [V1.3]

PRELIMINARY – SUBJECT TO CHANGE

Tensordyne Napier Pod

TDN72

The TDN72 is our foundational base unit. Our innovation in math, silicon, and scale-up networking is integrated into field-proven data center chassis. Air-cooled and drawing just 30kW to achieve maximum compatability with existing infrastructures.

System	TDN72 (Napier Gen)
Configuration	72x Tensordyne TDN AI Processors
AI Compute FP8 [1]	151.9 PFLOPS
AI Compute FP16 [1]	76.7 PFLOPS
HBM Capacity	10.4 TB @ 338 TB/s
Scale Up Bandwidth [2]	68.8 TB/s
Frontend IO Ethernet	8x dual-200-Gbps QSFP Ethernet
x86 cores	8x 10-core Intel Xeon D
Storage	64 TB
Physical Dimension	13RU (44.2 x 57.8 x 81.3 cm)
System Cooling	Air Cooled (Rear Fan Trays)
Operating Power	30 kW HVAC (180-305 VAC) 15A / 20A HVDC (190-410 VDC) 15A / 20A

[1] dense (no sparsity)
[2] total bidirectional system bandwidth (72 any-to-any direct connectivity)

Tensordyne Napier Artificial Intelligence Processor

TDN AIP

The TDN AIP is the world's first logarithmic math silicon for AI inference. By right-sizing the fundamentals of AI compute and combining with ultra-low latency fabric, it delivers massive throughput and efficiency. Together with an optimized blend of SRAM and HBM, it is ready for the most demanding workloads of today and well into the future.

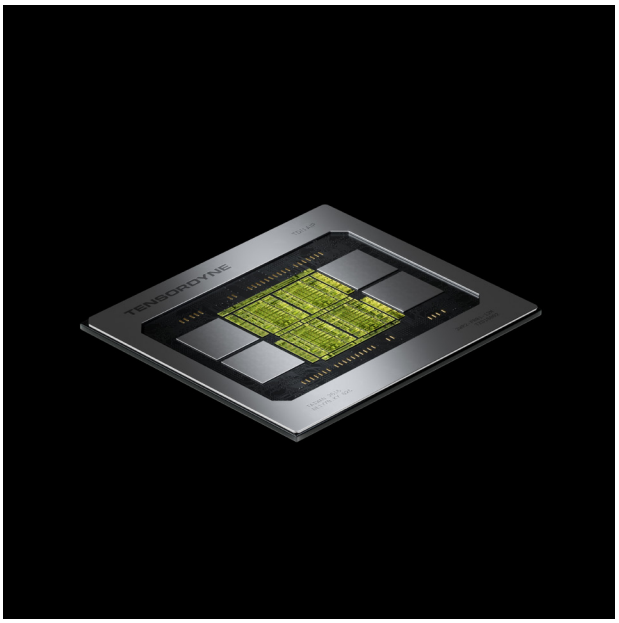

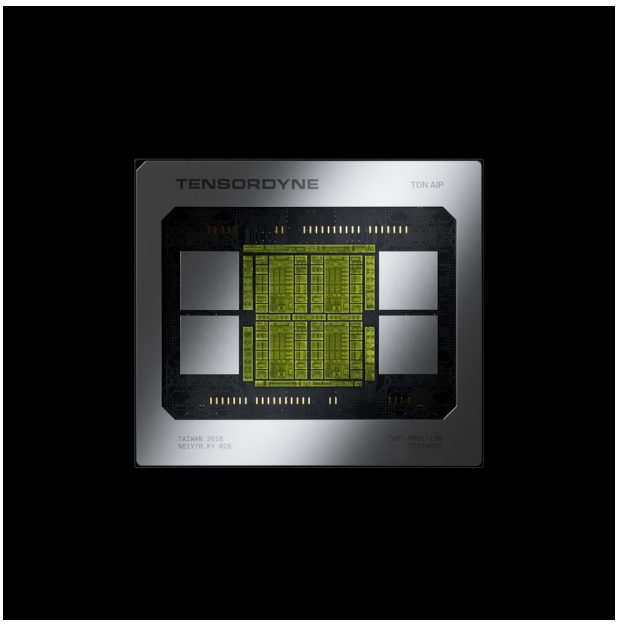

MCM	TDN AIP (Napier Gen) in 3nm TSMC
AI Compute FP8 [3]	2.11 PFLOPS
AI Compute FP16 [3]	1.07 PFLOPS
Vector Processing	Dedicated VPUs
Supported Precisions [4,5]	FP16, FP8, FP4, NVFP4
On-Chip SRAM Memory	256 MB @ 40 TB/s
HBM3E Capacity	144 GB
HBM3E Bandwidth	4.7 TB/s
Scale Up Fabric Latency	~ 1000 ns
Scale Up Fabric Bandwidth	1.0 TB/s
Host Connectivity	PCIe Gen5 x8
TDP	300 W

[3] dense (no sparsity)
[4] Dynamic scaling for all datatypes
[5] FP4 and NVFP4 supported only for parameters/weights

Tensordyne Napier AI Compute Tray

TDN ACT

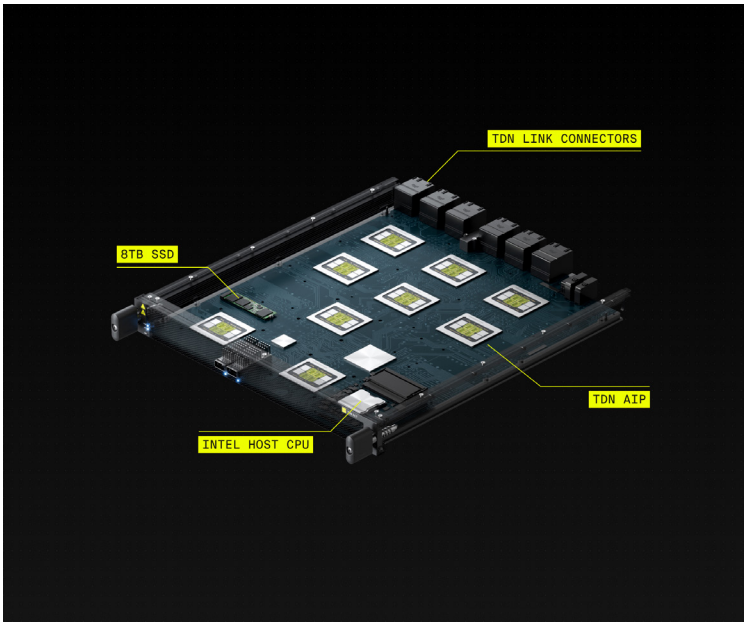

The TDN ACT is a marvel of modular engineering, pairing 9x TDN AIPs with a dedicated Intel Xeon host and 8TB of high-speed storage. It stays cool under pressure with a field-proven air-cooled design and slides seamlessly into the TDN Link backplane, delivering a connection that is as reliable as it is fast.

Tensordyne Napier Link

TDN LINK

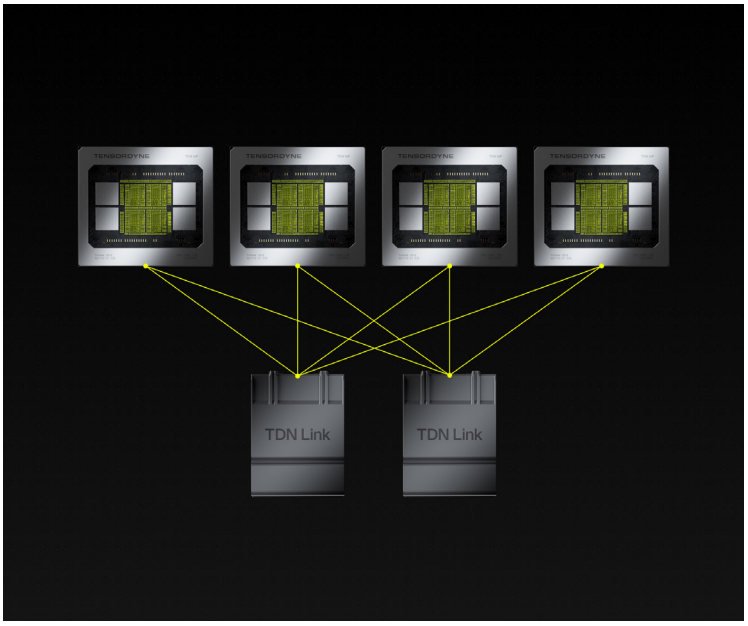

The TDN Link turns separate chips into a single, unified mind. By allowing 72 accelerators to work as one, this any-to-any, ultra-low latency fabric tackles the world's largest models with effortless speed and a fraction of the cost. It is the most advanced scale-up networking technology of its kind, delivering larger models and faster tokens without compromise.

Tensordyne Napier Rack

The TDN Rack is the ultimate expression of our inference-only architecture. Standing 52RU tall, it seamlessly integrates four TDN72 systems into a single AI powerhouse. It is built for total versatility, capable of running in standard high-efficiency token serving mode or scaling instantly through software-defined disaggregation. At just 120kW for the full rack, it delivers multi-trillion parameter performance with an elegance that is as efficient as it is powerful.

System	4x TDN72 (Napier Gen)
Configuration	288 Tensordyne TDN AI Processors
AI Compute FP8 [1]	608 PFLOPS
AI Compute FP16 [1]	307 PFLOPS
HBM Capacity	41 TB
Scale Up Bandwidth [2]	275 TB/s
Frontend IO Ethernet	32x dual-200-Gbps QSFP Ethernet

[1] dense (no sparsity)
[2] total bidirectional system bandwidth (72 any-to-any direct connectivity)

Highlighted Use Cases

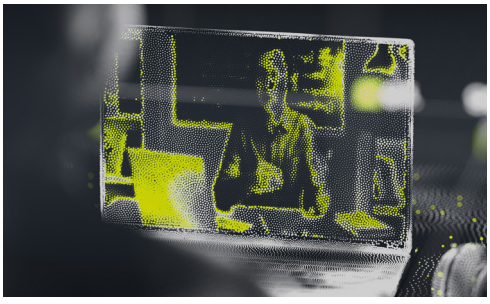

Wan2.2 AI Video

→ 4K resolution in 30 FPS realtime
→ $1.00 / 60 sec clip

DeepSeek R1 / Kimi K2.6

→ 1,000 OTPS per user
→ $0.30 / 1M Tokens

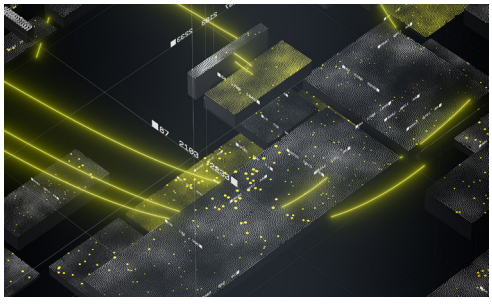

Agentic AI

→ Agentic task latency from seconds down to milliseconds

Tensordyne SDK

- → Huggingface Model Hub
- → PyTorch and Triton Model Definition Support
- → td.nn and td.ir exposure in Python for combined high-level layer definition and low-level kernel optimization
- → AI agentic transpilation migrates GPU-targeted code to Tensordyne's PyTorch-based domain-specific language.

Reach Out For More Information.

WWW.TENSORDYNE.AI
CONTACT@TENSORDYNE.AI

UNITED STATES
SUNNYVALE, CA 94089
1195 BORDEAUX DR

GERMANY
MUNICH, 80333
BRIENNER STRASSE 59

GILLES BACKHUS
CO-FOUNDER

GILLES.BACKHUS@TENSORDYNE.AI

NICK VITRO
VP SALES & BUSINESS DEVELOPMENT

NICK.VITRO@TENSORDYNE.AI

TENSORDYNE