


PRODUCT BRIEF: AMD PENSANDO™ VULCANO 800 AI NIC



OVERVIEW


The AMD Pensando™ Vulcano 800 AI NIC enables high-performance scale-out GPU networking and front-end host networking for AI workloads with Ethernet connectivity up to 800 Gbps per NIC and is the only NIC on the market to offer up to 2.4 Tbps of scale-out bandwidth per GPU.

Built on the proven AMD Pensando™ P4 programmable architecture, the AI NIC is fully hardware and software programmable, supporting UEC-ready RDMA, Multipath Reliable Connection (MRC), and other intelligent custom network transports.


That same programmability lets customers keep innovating at the speed of software. Scaled out across AI clusters, the AI NIC delivers fast AI job completion times, high resiliency, efficient link utilization, and excellent load distribution.



VULCANO PCIe CARD



VULCANO OCP 3.0 CARD



VULCANO AI NIC BOARD

KEY BENEFITS

UP TO 13% FASTER AI JOB COMPLETION TIMES¹

Accelerate AI job completion times to enable improved performance for demanding AI applications.

BOOSTED CLUSTER UPTIME

Help improve reliability, availability, and serviceability (RAS) to reduce interruptions in large-scale AI workloads.

UP TO 33% LOWER SWITCHING COSTS²


Optimize network switching costs by scaling AI clusters on fabrics that need fewer transceivers, cables, and optics.

ENHANCED OPERATIONAL EXCELLENCE

Roll out mission critical upgrades without sacrificing cluster uptime to maintain production readiness.

The AI NIC is available in multiple form factors: PCIe HHL, OCP TSFF, and 4-NIC and 6-NIC board form factors for the AMD Helios rack-scale platform.

PRODUCT BRIEF: AMD PENSANDO™ VULCANO 800 AI NIC




SPECIFICATIONS

	PCIe	OCP	Vulcano AI NIC Board - 6	Vulcano AI NIC Board - 4
TOTAL BANDWIDTH (UP TO)	800 Gbps		4.8 Tbps	3.2 Tbps
BANDWIDTH PER GPU (UP TO)	800 Gbps		2.4 Tbps	1.6 Tbps
FORM FACTOR	HHHL PCIe CEM	OCP 3.0 TSFF	Custom	
CONNECTOR	1xOSFP-800		3x1.6T OSFP	2x1.6T OSFP
ETHERNET INTERFACE	1x800G / 2x400G / 8x100G		6x800G / 12x400G / 24x200G	4x800G / 8x400G / 16x200G
NIC COOLING	Air-Cooled		Liquid-Cooled	
MANAGEMENT	PLDM over MCTP/USB RBT NC-SI Over MCTP/SMBus			
BOARDS PER COMPUTE TRAY	1	8/4	2	

INTELLIGENT TRANSPORT MODE OPTIONS

	AI NIC-Based Packet Spray	AI NIC-Based Source Routing	Switch-Based Packet Spray
NETWORK INTELLIGENCE LAYER	AI NIC	AI NIC	Switch
BEST SUITED FOR	Standard ECMP fabrics where switch simplicity is a priority	Multi-plane deployments require deterministic, operator-controlled routing.	Environments where switch-layer traffic control is the operational preference.

PRODUCT BRIEF: AMD PENSANDO™ VULCANO 800 AI NIC




FEATURES

AMD Pensando™ Vulcano 800 AI NIC		
SUPPORTED TRANSPORT PROTOCOLS	Standard Transports RoCEv2 (RDMA)	Custom Transports - UEC-ready RDMA - Multipath Reliable Connection (MRC) - Other Custom Transports
UEC-READY RDMA AND MRC PROGRAMMABLE FEATURES	- Intelligent Packet Spray - Out-of-order Packet Handling and In-order Message Delivery	- Selective Retransmission - Path Aware Congestion Avoidance - Rapid Fault Detection
CUSTOM TRANSPORT MODES	- NIC-based Packet Spray - NIC-based Source Routing (SR / SRv6)	Switch-based Packet Spray
SOFTWARE AND DRIVERS	- Standard Linux Drivers - DPDK - XDP, AF-XDP	
STORAGE NETWORKING	- RDMA and RoCEv2 - AI NIC Direct RDMA for GPU - AI NIC Direct RDMA for Storage	

DEPLOYMENT OPTIONS

STANDARD ADAPTER FORM FACTORS		BOARD FORM FACTORS FOR HELIOS RACK	
To support mainstream GPU servers and OCP-compliant AI platforms, the AMD Pensando™ Vulcano 800 AI NIC is available as standard adapter cards using a PCIe® host interface, delivering up to 800 Gbps connectivity for scale-out AI networking.		To support AMD Instinct™ GPUs in Helios rack based systems, the AMD Pensando Vulcano 800 AI NIC can be deployed with four AI NICs (two NICs per GPU) or six AI NICs (three NICs per GPU) to provide up to 2.4 Tbps of scale-out bandwidth per GPU.	
PCIe HHHL Adapter	OCP 3.0-TSFF Adapter	Vulcano AI NIC Board - 4	Vulcano AI NIC Board - 6

PRODUCT BRIEF: AMD PENSANDO™ VULCANO 800 AI NIC



VULCANO AI NIC ORDERABLE PART NUMBERS (OPN)

PCIe – HHL*	100-700000003
OCP-3.0 TSFF*	100-700000004
Vulcano AI NIC Board - 4	100-700000015
Vulcano AI NIC Board - 6	100-700000016

*All PCIe and OCP card adapters are shipped with the tall bracket mounted and a short bracket as an accessory.

DISCLAIMERS


The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED “AS IS.” AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

COPYRIGHT NOTICE

AMD, the AMD Arrow logo, Pensando and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe® is a trademark of PCI-SIG Corporation. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. Certain AMD technologies may require third-party enablement or activation. Supported features may vary by operating system. Please confirm with the system manufacturer for specific features. No technology or product can be completely secure. © 2026 Advanced Micro Devices, Inc. All Rights Reserved.

PRODUCT BRIEF: AMD PENSANDO™ VULCANO 800 AI NIC



ENDNOTES

1. Based on AMD Engineering silicon modeling and AMD synthetic benchmark simulation, using the MOE_4p5T hi_sparsity_ea9 benchmark test to project time- to- solution speed up in days using the FP8 training datatype on a simulated LLM system modeled with 8000 AMD Instinct MI455X GPUs and (3) versus (2) AMD Pensando Vulcano NICs. Results reflect analysis of pure training compute (decode layers only) Configuration evaluated with a global batch size of 4096 and a sequence length of 4K, using SwiGLU activation. Results exclude evaluation and checkpointing overhead. FlashAttention v3 with matmul–softmax overlap is assumed. AllReduce and All2All communication costs are fully accounted for and not hidden via tiled compute–communication overlap. Gradient synchronization and FSDP weight prefetch are evaluated across varying levels of overlap to assess scale-out sensitivity. Results assume ideal and not fully optimized real-world behavior and may vary when actual product(s) are released in market. MI400-018
2. PEN-022: AMD comparison and pricing as of May 18, 2026, for network fabric costs to support 32,000 GPUs. Comparison of a Vulcano-based NIC (VULCANO-CUSTOM-2.4T) deployed as part of a Helios rackscale system with a network based on 1.6T Tomahawk 6 switching with 200G SerDes versus using a competitor 800G NIC with 800G Tomahawk 6 switching with 100G SerDes. Both fabrics were fat-tree topologies built on Tomahawk 5 800G switching platforms, with NIC costs considered comparable. The Vulcano-based design is estimated to deliver up to 33% savings in network switching costs by enabling a more cost-effective architecture with fewer switching platforms, more bandwidth per port on the network, and reduced transceiver cables/optics.

TH6-100G Serdes Fat-Tree (Competition):

Switching (TH6C BCM78914 - 128x800G):

• Leaf Units	1,000
• Spine Units	500
• Total Switches	1,500
• TH6 100G Unit Price	\$79,587
• Total Switching Cost	\$60M

Cables/Optics:

• NIC Transceivers (800G-DR8)	64,000 @ \$500
• Leaf/Spine Transceivers (800G-DR8)	192,000 @ \$500
• MPO Cables	256,000 @ \$89
• Total Optics/Cables	\$64M

Total Fabric Cost TH6-100G (Switches + Cables/Optics): \$124M

TH6-200G Serdes Fat-Tree (Vulcano-Custom2.4T / AMD Solution):

Switching (TH6P BCM78910 - 64x1.6T):

• Leaf Units	1,000
• Spine Units	500
• Total Switches	1,500
• TH6 200G Unit Price	\$66,336
• Total Switching Cost	\$50M

Cables/Optics:

NIC Transceivers (1.6T-DR8)	32,000 @ \$900 (50% fewer vs. competition)
Leaf/Spine Transceivers (1.6T-DR8)	64,000 @ \$900
MPO Cables	128,000 @ \$89 (50% fewer vs. competition)
Total Optics/Cables	\$43M
Total Fabric Cost TH6-200G (Switches + Cables/Optics):	\$93M

Capex Savings (Fabric only):

Savings \$:	\$30.7M
Savings %:	33.1%

Pricing sources: SemiAnalysis Hyperscaler Networking Model data used with permission; full analysis available via SemiAnalysis subscription. Edgecore switch pricing as of May 17, 2026. Results may vary based on system configuration.