

AMD POLLARA 400 CARD



OVERVIEW

The AMD Pollara 400 delivers a fully programmable 400 Gigabit per second (Gbps) RDMA Ethernet Network Interface Card (NIC).

The Pollara 400 PCIe NIC builds upon the success of the proven AMD Pensando P4 architecture by combining a high-bandwidth Ethernet controller with a unique set of highly optimized hardware acceleration engines to enhance network performance and improve AI job completion times.


ADVANCING RDMA NETWORKING

AMD Pollara 400 is an advanced solution designed to optimize backend networking using a fully programmable Remote Direct Memory Access (RDMA) transport and hardware based congestion control. Pollara 400 helps facilitate efficient data exchange by passing components not required for GPU-to-GPU communication, reducing latency and increasing throughput. AMD enhanced transport is highly scalable, delivers low latency, can run on any Ethernet fabric, and offers some key innovations for AI:

- Intelligent Packet Spray
- In-Order-Delivery (messages to GPU)
- Selective Retransmission
- Path Aware Congestion Avoidance

The Pollara 400 seamlessly integrates into standard compute servers, delivering high-performance networking specifically optimized for AI and ML workloads while reducing complexity at scale. Pollara 400 is UEC ready, offers RoCEv2 compatibility, inter-operability with other NICs, making it a valuable addition for training and inference use cases.

AMD POLLARA 400 PCIe® CARD



SPECIFICATIONS

MAX BANDWIDTH	• 400 Gbps
FORM FACTOR	• Half-height, half-length (HHHL)
HOST INTERFACE	• PCIe Gen5.0 x16
ETHERNET INTERFACE	• QSFP12 (NRZ/PAM4 Serdes)
ETHERNET SPEEDS	• 25/50/100/200/400 Gbps
ETHERNET CONFIGURATIONS	• Supports up to 4 ports - 1 x 400G - 2 x 200G - 4 x 100G - 4 x 50G - 4 x 25G
MANAGEMENT	• MCTP over SMBus

ACCELERATED AI NETWORKING

Purpose-built for hyperscale AI workloads, the UEC ready programmable Pollara 400 combines high bandwidth, low latency Ethernet controllers with a specialized set of hardware acceleration engines designed to accelerate GPU-to-GPU communication, helping ensure high-bandwidth data transfers and improve AI job completion performance with low power consumption.

The adapter supports standards-based RDMA over Converged Ethernet (RoCE), the unique fully programmable AMD RoCE transport and congestion control pipeline, provides the flexibility needed by the largest cloud customers to implement their own advanced congestion control algorithms.

These capabilities are designed to optimize AI backend networks bypassing hotspots, increasing throughput and facilitating efficient data exchange for training jobs. AMD Pollara 400 can provide the foundation as customers move from proprietary technologies to open, standards based Ethernet for their large-scale backend training datacenters.

WHY AMD

Generative AI is rapidly reshaping the IT landscape, with AMD at the forefront through its pioneering innovations in data centers and networking. One significant advancement is the Pollara 400, a network accelerator designed to optimize AI workloads over Ethernet-based networks. Ethernet enables high-speed GPU-to-GPU communication, with the Pollara 400 achieving network performance of up to 400Gb/s with Remote Direct Memory Access (RDMA) over Converged Ethernet.

As AI advances, fueled by generative AI and large language models, it is unlocking new capabilities and enabling computers to perform tasks with greater demand on the infrastructure. The success of AI is increasingly reliant on GPU-accelerated computing for processing large datasets, training extensive AI models, and facilitating real-time inference. However, this shift presents challenges for Ethernet networks. Originally built for compatibility and tiered applications, Ethernet now needs to address the demands of modern AI workloads, which require tightly coupled parallel processing, rapid data transfers, and specialized communication patterns—necessitating optimized network connectivity.

The new AMD transport protocol aims to enhance the performance of AI/ML workloads by optimizing congestion control and load balancing. It features an adaptive load balancing algorithm, multi-bit congestion indicators, out-of-order delivery handling, selective retransmission, and swift loss recovery—all integrated into hardware. The Pollara 400 is a state-of-the-art solution designed to optimize backend networking through a fully programmable RDMA transport and hardware-based congestion control. This programmable transport and congestion control system allows for precise congestion management, helping ensure the fastest completion times for AI workloads. AMD innovation is evident in its data path-driven congestion control, which is integrated into the programmable transport rather than relying on a separate CPU, as seen in other architectures. These advancements can significantly enhance data center performance, efficiency, security, and help reduce total cost of ownership (TCO).

DEPLOYMENT OPTIONS

AI COMPUTE NODE

Typical deployment of Pollara 400 is targeted at back-end networks enabling 400G GPU communication acceleration for massively scalable AI clusters while enabling deterministic latency and high performance. Pollara 400's can be installed in every AI compute node in the data center supercharging hyperscale AI workloads. Pollara 400 supports leadership remote direct-memory access over Ethernet (lossy or lossless) delivering optimal AI workload efficiency. Pollara 400 software interacts with collective communication libraries to optimize job completion times at the lowest latency possible.

SOFTWARE FUNCTIONALITY

KEY FEATURES*

ENHANCED OBSERVABILITY

- Latency metrics, drop statistics,
- Telemetry threshold alerting
- Inline Hardware IPsec encryption and decryption
- AES-GCM 128/256-bit key

STORAGE ACCELERATION

- RDMA
- RoCEv2
- UEC Ready RDMA

MANAGEMENT AND CONTROL

- MCTP/SMBus
- SPDM over MCTP
- MCTP over PCIe VDM
- PLDM firmware

CLOUD NETWORKING

- UEC ready
- Flexible Transport
- Multi-Tenancy

OPTIMIZED AI NETWORKING

- Programmable Congestion Control
- GPUDirect
- Path-Aware Congestion Avoidance
- Intelligent Packet Spray
- In-Order-Delivery & Loss Retransmission

ADVANCED SECURITY

ADAPTER CARD PORTFOLIO AND FORM FACTORS

ORDERABLE PART NUMBER (OPN)	• Pollara-400-1Q400P
HOST INTERFACE (PCIe)	• PCIe Gen5.0 x16
FORM FACTOR	• HHHL
BANDWIDTH	• 400Gb/s
NETWORK PORTS	• 1 x QSFP112

All adapters are shipped with the tall bracket mounted and a short bracket as an accessory

DISCLAIMERS (FN -1)

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 (FN-1)

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-part 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. © 2023-2024 Advanced Micro Devices, Inc. All Rights Reserved.