

AMD PENSANDO: SILICON THAT SUPERCHARGES THE MODERN DATACENTER

EXECUTIVE SUMMARY

Enterprises face challenges on multiple fronts, with new demands on datacenter infrastructure. Emerging workloads, born of today’s generative AI “gold rush,” are computationally intensive and tax the limits of legacy hardware and software. Edge enablement, designed to lower latency and improve network performance, also increases power consumption and corresponding operational expense levels. The latter is problematic, given that most enterprises are engaged in sustainability efforts to reduce carbon emissions and curb power consumption. Furthermore, the rise of modern work-from-anywhere management creates management complexity with the need to manage connectivity and security at scale across domains — from on-premises, to edge, to multi-cloud.

To meet these requirements, modern datacenters must undergo a rapid transformation. Subsequently, they require high availability to address the non-stop operational needs of organizations of all sizes. Scalability is also critical to meet growing workload and distributed application demands. Finally, hardened security must be tightly woven into the overall connectivity stack to ensure resiliency and efficacy.

Data processing units (DPUs) that integrate into underlying IT and OT infrastructure address these needs. While a handful of purpose-built silicon solutions are available, organizations require guidance to choose the right platform. Moor Insights & Strategy believes that AMD Pensando is well positioned to deliver the ideal DPU platform given its market maturity, accompanying full software stack, backward compatibility, emerging workload support, and hardened security.

THE VALUE OF DPUS

What is a DPU? Its core architectural design is programmable and infuses general-purpose CPU functionality with network interface capabilities. DPUs provide compelling functionality in offloading workload overhead from server CPUs to boost performance for specific applications. Through network traffic offload, DPUs can increase infrastructure operational efficiency, accelerate performance, reduce power to meet sustainability goals and provide higher scalability, availability, and security.

The value of a DPU's capability cannot be understated. Enterprises that deploy DPUs can dramatically mitigate operational expenses through less power consumption and realize longer term investment protection in deployed infrastructure given their programmable nature. A DPU is purpose built for data processing within datacenter environments. It marries a rich set of acceleration engines to multi-core CPUs that can improve the performance of AI and machine learning algorithms and facilitate zero-trust deployments with enhanced encryption and silicon root of trust. However, not all DPUs are equal. Moor Insights & Strategy believes that AMD Pensando stands out with a complete solution.

WHY AMD PENSANDO

AMD Pensando is a fully programmable DPU designed to deliver cloud, compute, network, storage, and security services at scale. Its deep capabilities are used by large public cloud service providers with massive install bases, such as Microsoft Azure and Oracle Cloud Infrastructure, which require the highest levels of scalability, deployment flexibility, performance, security, and investment protection. However, AMD Pensando is also well suited for smaller datacenter deployments due to its software stack, backward and forward compatibility, and consolidated functionality.

First, AMD Pensando provides the industry's only bundled software stack. This contrasts with competitive DPU offerings that only provide software development kits (SDKs) to facilitate developer programmability. Slower integration into datacenter infrastructure often occurs in these scenarios, despite the Linux Foundation's significant efforts to standardize application program interfaces (APIs) through the creation of the underlying software stacks based on open standards and broad ecosystems. AMD Pensando simplifies this effort and, in the process, speeds deployment and ensures consistency. AMD also supports developer efforts with AMD Pensando through its software-in-silicon development kit (SSDK). The result is an extremely flexible path for organizations that want to embrace the power of DPUs.

Second, the AMD Pensando DPU platform provides backward and forward compatibility, delivering exceptional investment protection. In its third generation with the 2x400G Salina iteration, AMD offers one of the most market-mature DPU platforms in the industry. By ensuring backward compatibility, AMD provides datacenter operations professionals smoother, uninterrupted performance with the latest features, higher throughput, and lower latency. Given the rise of generative AI and its potential to tax IT infrastructure from a computational and networking standpoint, AMD Pensando delivers the headroom needed for future datacenter workloads and applications.

Finally, AMD Pensando DPUs consolidate routing, switching, security, and storage functionality. The momentum towards shrinking datacenter footprints and the corresponding power consumed by IT infrastructure makes this a significant feature. As a result, AMD Pensando delivers a consolidated, top-of-rack solution that can manage multiple functions from a single management console. AMD's success is evidenced through its development efforts with large hyperscalers in addition to Hewlett Packard Enterprise (HPE) Aruba Networking and others. For enterprises considering the use of DPUs to streamline critical network, storage, and security functions, Moor Insights & Strategy believes that AMD Pensando can dramatically mitigate capital and operational expenses tied to datacenter operations and do so in a highly efficient manner.

AMD PENSANDO FAMILY

Diving deeper technically, AMD Pensando offers a mature family of infrastructure accelerators evidenced by its multi-generational portfolio. Each provides full programmability that utilizes industry-standard P4 language to support many solutions. This design allows AMD Pensando to execute an optimized software stack that delivers cloud, compute, network, storage, and security services at scale. For flexibility, AMD Pensando can also be deployed in several configurations, including PCIe cards, network and security appliances, and SmartSwitches

Furthermore, AMD Pensando can deliver these mission-critical enterprise services with an exceptional latency profile that provides high performance while operating at a low power-consumption envelope. Many enterprises are committed to meeting aggressive environmental, social, and corporate governance (ESG) objectives, and the use of AMD Pensando can help in these efforts. Furthermore, with an advanced architectural design and feature set, AMD Pensando can support a host of datacenter applications that provide exceptional connectivity, security, and observability functionality, including:

- Advanced networking functions, including full software defined networking (SDN) support and load balancing.
- Security features that span stateful firewall, stateful connection tracking, security groupings, IPsec and TLS/ DTLS encryption.
- Storage offloads for both initiator and target (NVMe/TCP).
- Robust observability support in the form of flow-based packet telemetry and stateful connection tracking.

- Programmability that provides headroom for next-generation, generative AI and machine learning workload training and inferencing, leading towards standardization within the Ultra Ethernet Consortium.

DXC TECHNOLOGY LEVERAGES AMD PENSANDO FOR ITS MODERN DATACENTER OPERATIONS

DXC Technology is an IT services company founded over a half-century ago and operating in 60 countries. It delivers engineering, IT outsourcing, cloud, modern application, and security solutions to nearly half the Fortune 500. Consequently, DPUs represent an opportunity for the company to modernize its datacenter infrastructure stack and offer more resilient, high-performance services to its customers at cloud scale.

DXC Technology turned to HPE Aruba Networking when it recently embarked on a reimagining effort tied to its datacenter footprint consolidation. The company's goal was to move from approximately 200 datacenters to 60. HPE was a logical choice, given it was the first company to embrace the AMD Pensando DPU in its [Aruba Networking CX 10000 Series Switch](#). As a result, DXC Technology claims that it has improved operational efficiency by up to 50%, while realizing new-found agility, visibility, and the ability to meet its sustainability objectives. The company has also measured a significant total cost of ownership savings of \$65 million through the deployment of the HPE Aruba Networking CX 10000 platform.

Ongoing co-development efforts with HPE and AMD Pensando enable DXC Technology to refine its routing, switching, and firewall security operations. Today, among many other benefits, DXC Technology leverages the power of the underlying infrastructure for east-west traffic segmentation, and future north-south capabilities will go far to improve visibility and performance given Pensando's highly programmable nature.

CALL TO ACTION

Enterprises are facing challenges that are stretching the limits of existing datacenter infrastructure. Generative AI and edge enablement are poised to supercharge use cases, but power-hungry new workloads and the need to manage connectivity and security at scale across domains create management complexity. DPUs have the potential to address these needs. With market maturity, accompanying a full software stack, backward compatibility, emerging workload support, and hardened security, Moor

Insights & Strategy believes that AMD Pensando is well positioned to deliver the ideal DPU platform.

To learn more, [visit AMD Pensando Infrastructure Accelerators](#).

CONTRIBUTOR

[Will Townsend](#), Vice President & Principal Analyst, Networking & Security Practices at [Moor Insights & Strategy](#)

PUBLISHER

[Patrick Moorhead](#), Founder, President, & Chief Analyst at [Moor Insights & Strategy](#)

INQUIRIES

[Contact us](#) if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

AMD commissioned this paper. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2024 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.