Enterprise Strategy Group

SHOWCASE

JUNE 2025

The Critical Role of NIC Programmability in Scaling Out Data Center Networks for AI

Jim Frey, Principal Analyst, Enterprise Networking

Abstract: Infrastructure buildouts are underway for hosting AI workloads. Networks play a critical role in effective scale-out, and Ethernet is expected to become the majority choice in these networks this year. But effective networking isn't just about switches; building advanced functionality into network interface cards is an essential design strategy. AMD Pensando offers advanced, programmable network interface cards (NICs) that represent an optimized path to success.

Building Infrastructure for AI

The race for AI is off and running, with organizations of all sizes working to apply generative AI (GenAI) and agentic AI to business problems and work processes in the service of greater efficiencies and improved business advantage. But these AI workloads are unique when compared to most other demands on IT infrastructure, requiring high power, high efficiency, and high throughput compute, storage, and networking. While many teams started working with GenAI using cloud-based services, longer-term plans for extended usage must account for cost, data privacy, and governance constraints, leading most enterprises to plan building GenAI-specific infrastructure closer to home. In fact, a recent study by Enterprise Strategy Group found that 89% of organizations have built and deployed their own custom GenAI infrastructure on premises or will do so within the next 24 months.[1]

Building on-premises infrastructure for GenAI and agentic AI carries very specific technical requirements. While GPU-based computing has been most widely publicized as being at the core of AI infrastructure, there are other essential supporting components, including, most notably, high-performance storage and high-performance, non-blocking networking. Those networks connect the servers and the storage, both within AI infrastructure clusters (i.e., east-west or "scale-out" networking), as well as to the rest of the IT infrastructure outside of the clusters (i.e., north-south networking, through the frontend network).

Scale-out networks for AI infrastructure are commonly delivered using one of two networking technologies: InfiniBand or lossless Ethernet known as RDMA over Converged Ethernet (RoCE). While InfiniBand has found the most early usage among cloud providers, going forward, cloud providers and enterprises alike are favoring ethernet-based solutions as the networking standard for on-premises AI infrastructure buildouts. This is primarily due to the familiarity that networking teams have with Ethernet and the existing relationships with vendors providing RoCE products, among other reasons (see Figure 1).

[1] Source: Enterprise Strategy Group Research Report, *Architecting the Network for an AI-powered World*, February 2025. All Enterprise Strategy Group research references and charts in this showcase have been taken from this report.

Figure 1. Enterprises Favor Ethernet for AI Infrastructure Networking

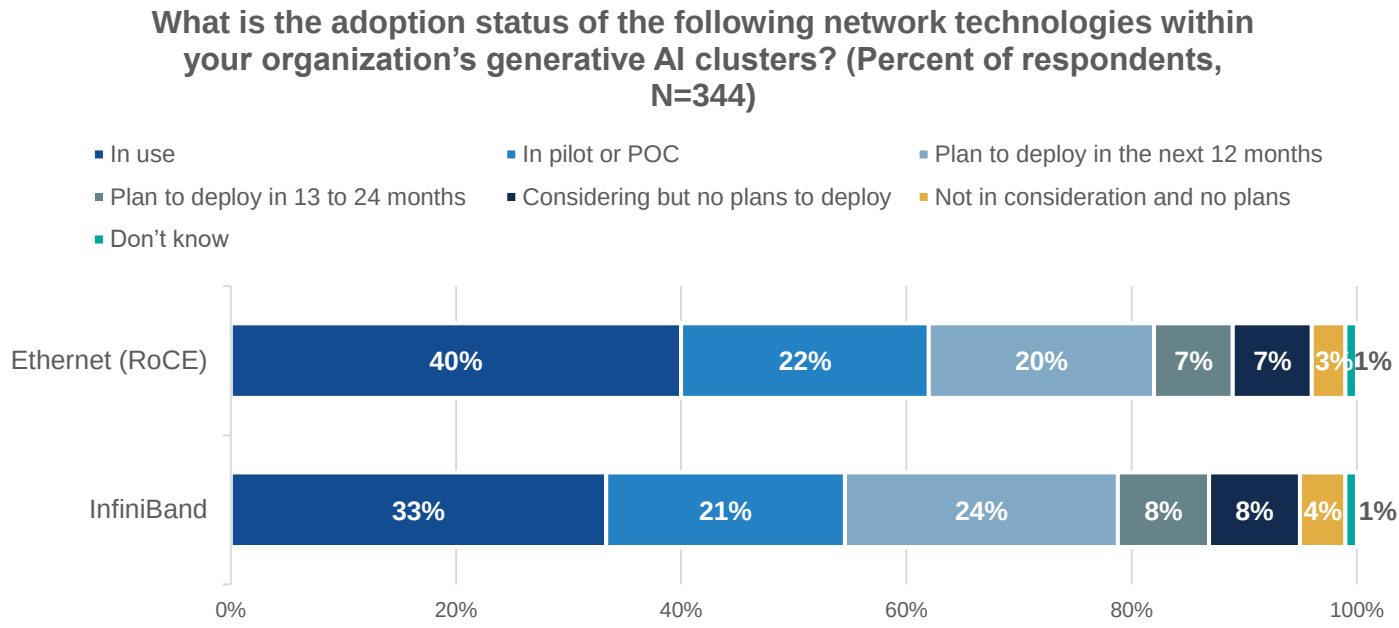

Source: Enterprise Strategy Group, now part of Omdia

While there are many network objectives for AI infrastructure clusters, high resilience and high performance are at the top of the list. Network performance demands in this setting are particularly intensive due to the cost of AI model training and the extreme sensitivity to efficient movement of data between cluster nodes (see Figure 2).

Figure 2. Top Priorities for AI Infrastructure Networking

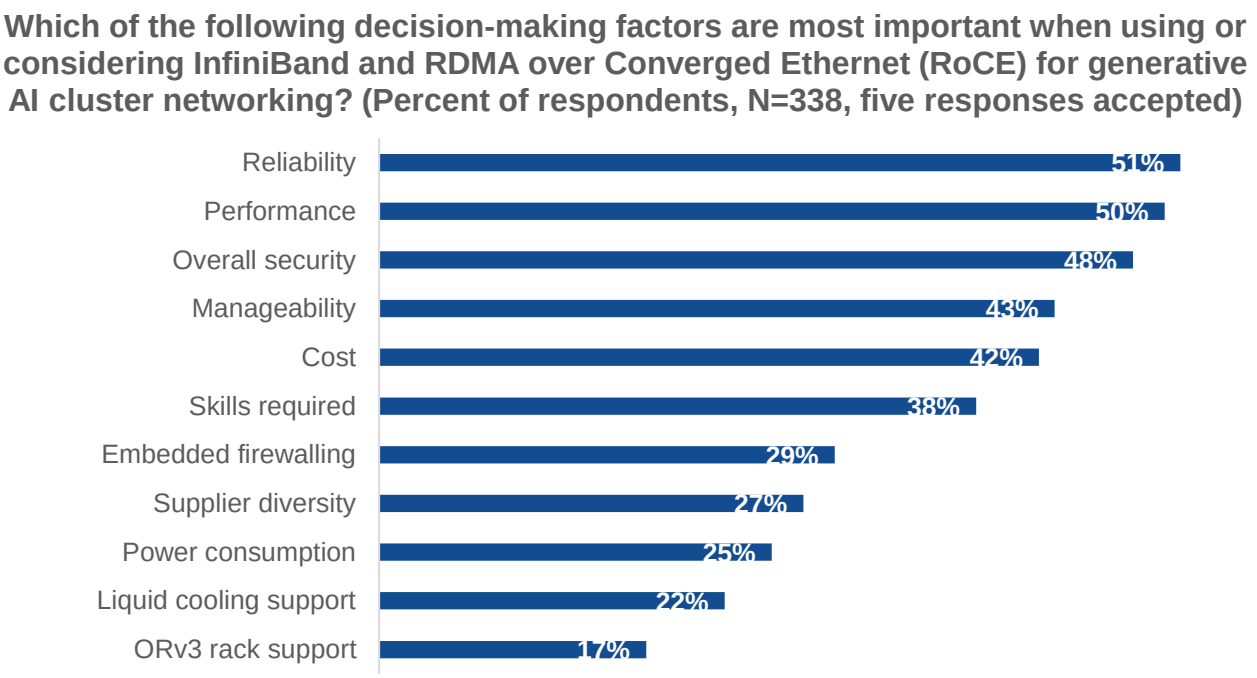

Source: Enterprise Strategy Group, now part of Omdia

The Role of Network Interface Cards in AI Infrastructure

In the case of RoCE, RDMA NICs provide direct memory access between GPU compute notes and other devices across the network, and just as with any complex IT infrastructure, the scale-out networks for AI infrastructure are not just about switches. The NICs in cluster servers are also an essential part of the design. NICs must go beyond their traditional role of network access to help deliver the high performance and high resilience demanded by the overall system.

Ethernet has traditionally been used to deliver a flexible, lossy network, but scale-out networks, especially those used for AI, need to ensure high throughout, low latency, and low jitter for optimal training performance. These new fabrics bring additional complexity that many organizations struggle with, hence NICs that can better handle load balancing and congestion at the end node can help enable scale-out fabrics to be somewhat lossy and less operationally rigid.

Further, and ideally, NICs that are based on chip technology that adds optimizations and programmability designed specifically for AI infrastructure applications offer a means to improve results and support future needs.

For enterprises building out AI infrastructures, the best NIC solutions will be designed to deliver a range of specific and highly valuable capabilities, including:

- **Features relevant to the unique needs of AI infrastructure networking.** Off-the-shelf commodity NICs will simply not work for AI infrastructure. NICs need to have been designed specifically to deliver features that optimize performance and flexibility to support overall infrastructure efficiency and efficacy.
- **Proven support for industry standards and features.** The Ultra Ethernet Consortium, or UEC, was founded in 2023 to establish a concensus for standards around the specialized adaptations included in RoCE for delivering high-performance, distributed, and lossless transport. The UEC's objectives are to lay the groundwork for interoperability across a range of vendors and establish RoCE as a viable alternative to InfiniBand by specifically optimizing it for the performance needs of large-scale AI workloads. Many Ethernet switch vendors support RoCE, and by adhering to UEC standards, NICs can ensure successful operations regardless of which switch vendor and which type of ethernet (RoCE or not) is chosen.
- **Programmability for tunability and future flexibility.** While standards are an essential starting point, each organization's AI infrastructure will also have unique characteristics based on the specific infrastructure components in use and the software used to train models, manage data, and execute inferencing. Consequently, every AI infrastructure instance must be tunable to adjust operations in ways that can reduce training times and eliminate wasted capacity.

AMD Solutions for Programmable, AI-ready Network Interfaces

The AMD Pensando™ Pollara 400 AI NIC is programmable, UEC-ready, and purpose-built to maximize performance in the largest AI scale-out clusters, accelerating both RDMA and RCCL (Radeon Collective Communication Library, an open-source library for optimizing multi-CPU environments) communications that underpin AI infrastructure operations. The Pensando Pollara 400 AI NIC enhances network efficiency with advanced congestion management and load balancing capabilities, delivering the high throughput, predictable performance, and scalability critical for next-generation AI infrastructure. Important capabilities include:

- 400 Gbps connectivity to AI infrastructure cluster networks supporting either lossless or lossy fabrics.
- Full support of UEC standards, enabling the solution to work with any UEC-compliant network switches.
- Open ecosystem to preserve customer choice and avoid vendor lock-in.
- Intelligent load balancing for optimizing traffic flows via capabilities such as adaptive packet spray.
- Multi-plane fabric support to localize faults, ensuring one plane failure doesn't impact others.

- NIC high-availability failover, eliminating single points of failure.
- Programmable transport, enabling rapid software innovations to deliver advanced UEC functionality and performance advantages without requiring hardware re-spins as standards evolve, unlike many other solutions. It also allows implementation of custom RDMA transport and congestion control protocols, optimized specifically for the needs of individual AI training cluster deployments.
- Active congestion management functions, including congestion avoidance, path-aware congestion control, and actional telemetry.
- Operational resilience via rapid fault detection, fast failover, and loss recovery.
- P4 programmability for tuning, flexibility, and adaptation to future needs.

The design richness and programmable flexibility of the AMD solution provides direct performance benefits in terms of faster job completion, improved network performance, and improved network reliability. It also helps on the business side of the equation in terms of lower total cost of ownership, open ecosystem support for avoiding vendor lock-in, and future readiness to adapt to new workloads and evolving standards.

Conclusion

AI adoption is occurring globally, and while cloud and SaaS will always play a role, organizations are actively planning to build out local AI infrastructure to meet business and security needs. With the emergence of the UEC and the open ecosystem growing around RoCE, Ethernet has emerged as the networking standard for these buildouts, which impacts both network switching and intelligent and high-functioning NICs installed in each cluster server. These NICs must be designed to support UEC standards and include specific features needed for these unique, high-cost, highly performant networking environments, all while offering programmability for tuning and adaptation. The AMD Pensando Pollara 400 AI NIC is an excellent example of a solution that has been built specifically for futureproof scale-out of AI infrastructures.

About Enterprise Strategy Group
Enterprise Strategy Group, now part of Omdia, provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

contact@esg-global.com
www.esg-global.com