

**EVIDEN**

**BullSequana  
eXascale  
Interconnect V3:  
Intelligent Network  
Management Accelerates  
GPU Performance in AI-HPC**



## Introduction

Throughout history, large-scale High-Performance Computing (HPC) has been instrumental to life-changing innovations. More recently, it has enabled disruptive technological advancements, laying the foundation for the rapidly emerging field of artificial intelligence.

As we **transition from the era of large-scale HPC to large-scale AI**, we unlock new opportunities and benefits for AI workloads. This shift comes at a time when the demand for more accurate predictions, such as in weather forecasting, and more performant AI models is pushing the boundaries of technological development at an unprecedented pace.

While the compute power of GPUs has grown significantly, following Moore's Law—which predicts that CPU power will double approximately every 18 months—networking infrastructure has not kept up at the same rate. Even though **bandwidth now doubles every two years**, currently reaching 400 gigabits per second and projected to hit 1600 gigabits per second in the near future, it remains insufficient to meet the increasing demands of modern compute environments. This gap between compute power and networking capabilities presents an urgent challenge.

To address this imbalance, **it is essential to accelerate networking advancements to keep pace with the rapid growth in computing power**. Closing this gap is vital for sustaining the progress of AI and other high-demand technologies. This paper explores the current state of networking and the necessary innovations to ensure it evolves in tandem with compute power, meeting the needs of the future.



## Why are existing high speed networking solutions inefficient?


Existing high-speed networking solutions are largely inefficient because they fail to address a critical bottleneck known as the «networking wall.» In large-scale AI training, it's often assumed that more GPUs are needed to handle increasing workloads, but in reality, the network itself is the limiting factor. While networking vendors continue to double bandwidth every few years, this alone is not enough to solve the problem. The issue lies in the fact that up to 70% of the time, GPUs are left idle, waiting for data due to networking delays.

The core of the problem isn't just about increasing bandwidth or «making the pipe bigger»; it's about managing the network more intelligently. Without a smarter networking solution, organizations are not fully utilizing their expensive hardware.

Instead of adding more GPUs to compensate for these delays, the solution is to optimize the network, allowing existing resources to perform at their full potential. By offloading the entire communication protocol to the NIC, we can increase the speed of the protocol being used and free up GPUs and CPUs from networking problems and latencies. This shift in thinking is essential for improving efficiency in large-scale computing environments. Smarter, not bigger, networks are key to unlocking the true power of modern AI infrastructure.

As GPU are waiting for the network 70% of the wall clock time for AI workloads, **Networking is the most important thing**





## The Ethernet Advantage

With the introduction of Eviden's **BXI V3**, we have chosen Ethernet as the standard protocol to broaden the accessibility of our technology, moving away from closed or proprietary systems that limit flexibility and interoperability. Ethernet, as the most widely used networking standard globally, offers the advantage of an open and widely compatible platform, enabling a larger user base to benefit from our advanced HPC and AI capabilities. By transitioning from a closed system to an open, Ethernet-based solution, we are eliminating the constraints of proprietary technology and opening the door to greater scalability and adaptability.

While Ethernet's open design fosters broad compatibility, it does not inherently meet all the demands of high-performance computing and AI workloads, such as addressing the «networking wall» and optimizing for large-scale operations.

To bridge this gap, Eviden has adapted the Ethernet protocol, reengineering it to better manage congestion, latency, and scale. This transition from a closed to an open system has allowed us to focus on reducing bottlenecks and improving latency management, unlocking the full potential of Ethernet for HPC and AI applications.

**BXI V3** leverages a specialized smart network interface card (NIC) to streamline communication, optimizing the entire protocol to free up GPUs and CPUs from the delays caused by networking inefficiencies. These enhancements allow us to take Ethernet to the next level, accelerating workloads and making it a more suitable choice for the high demands of HPC and AI environments. This approach helps us break through existing limitations, ensuring that Ethernet can support the future of large-scale computing.

# The Game-Changing Advantage of Eviden's BXI SmartNIC

For over a decade, Eviden has been refining Smart NIC technology, which lies at the core of **BXI V3**. The defining feature of BXI has always been the offloading of the communication protocol to the NIC. What makes our Smart NIC truly unique is its ability to not only send and receive data, but also process it, orchestrating communication across the network in ways that traditional NICs simply cannot. This offloading capability allows for asynchronous operations.


In a traditional, synchronous setup, GPUs remain idle while they wait for the NIC to complete data transfers. From a cost and efficiency perspective, a GPU that costs \$40,000 is only being used 30% of the time when networking inefficiencies exist, effectively increasing the cost per GPU to \$120,000. This means customers aren't getting the return on investment (ROI) they expect, as the true potential of their devices is limited by poor network performance.

With our Smart NIC, the GPU offloads its data to the NIC and immediately moves on to other tasks, eliminating costly idle time.

By reducing the **70%** wait time

typically experienced in synchronous setups, the Smart NIC frees up GPU resources, ensures GPUs are utilized more effectively, thereby significantly boosting ROI, reducing the total cost of ownership (TCO) and improving system performance and efficiency. This combination of advanced protocol offloading and seamless integration is a distinctive feature of BXI V3, offering a level of efficiency and performance that is unmatched in the market.

Additionally, **BXI V3** embeds the protocol directly into the NIC hardware, making it a plug-and-play solution. Applications run seamlessly without requiring changes, yet the performance is automatically enhanced. The Smart NIC also provides advanced congestion management, preventing network bottlenecks and ensuring higher throughput and lower latency, which leads to faster execution of applications.





# What makes BXI V3 truly unique

In addition to our unique Smart NIC, **BXI V3** provides several additional advanced features designed to enhance performance by reducing system overhead and streamlining data processing. For starters, **BXI V3** offers transparent virtual-to-physical address translation, allowing applications to directly post requests to the SmartNIC using virtual addresses without the need for system calls. This kernel bypass solution eliminates the need for memory registration and pinning, giving the kernel more flexibility to optimize memory allocation without interrupting the process. This improves efficiency by enabling seamless memory management while maintaining high performance.

**BXI V3** also enables registration of up to 32 million potential reception buffers, which the SmartNIC selects using matching keys based on message attributes. By offloading buffer matching to the SmartNIC, CPU resources are freed for computational tasks, reducing the need for CPU intervention in message processing and improving overall system efficiency. **BXI V3** also allows the registration of algorithms to be executed upon message reception, enabling the NIC to handle complex collective operations, including mathematical atomics. By offloading these tasks to the NIC, execution steps can be performed closer to the network, reducing overall execution time and further minimizing CPU involvement. This leads to faster processing of collective operations and more efficient network communications.

## Conclusion

As **technological demands continue to grow and evolve**, we are uniquely positioned to meet these challenges. As one of only four companies worldwide capable of delivering high-speed, low-latency networks—and the sole European provider—we stand apart. Our role as the only non-American founding member of the UEC further strengthens our ability to offer sovereign technology solutions, ensuring that **we remain at the forefront of the industry**. In addition to our work with the UEC, we are partnering with like-minded industry leaders, such as AMD, to drive innovation and shape the future of AI and high-performance computing.





EVIDEN

AMD 

Connect with us



[eviden.com](https://eviden.com)

Eviden is a registered trademark © Copyright 2024, Eviden SAS - All rights reserved.