DRIVENETS

Optimal AI Networking Fabric

- 10%-30% average JCT improvement
- Predictable, low latency ensures highest GPU utilization

Trusted Solution

- In Production cluster of leading Hyperscalers and Enterprise customers
- Running AT&T's IP core network since 2017
- Based on Standard Ethernet

DriveNets Network Cloud-AI

Offering a lossless, predictable network fabric that enables maximum utilization of GPUs, this is reflected in JCT that is faster than InfiniBand and any standard Ethernet solution

DriveNets is a rapidly growing company that has introduced a revolutionary approach to building mid to large-scale AI/HPC networking infrastructures. Traditionally, most AI/HPC networking infrastructures were hardware-centric, relying on monolithic systems and often proprietary technologies like InfiniBand. In contrast, DriveNets Network Cloud-AI is a software-based networking solution built on scheduled Ethernet fabric and a cloud-native architecture.

By combining the scalability of standard Ethernet with the high-performance of a chassis, DriveNets Network Cloud-AI delivers unmatched networking performance at any scale at an optimal cost structure.

DriveNets AI/HPC Scheduled Ethernet Fabric

DriveNets Network Cloud-AI

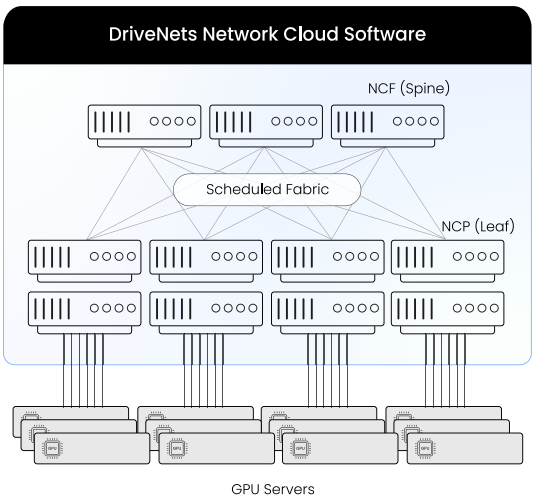

Large AI/HPC systems handle extensive workloads running across individual compute and storage nodes, which function together as a logical cluster connected through a high-capacity interconnect fabric, commonly referred to as backend AI/HPC network fabric. These networking solutions must continually evolve to maximize the utilization of expensive AI compute resources while supporting standardized connectivity to ensure vendor interoperability.

DriveNets Network Cloud-AI delivers industry-leading AI networking fabric performance at scale, supporting up to 32,000 GPUs in a single cluster. Built on scheduled Ethernet fabric architecture, it provides predictable, lossless back-end cluster connectivity with a proven 10%-30% improvement in job completion time (JCT) for high-performance AI workloads.

With a GPU, ASIC, and ODM-agnostic design, DriveNets Network Cloud-AI ensures flexibility and adaptability while delivering over 100% return on investment (ROI) from day one. This solution sets a new standard for scalability, efficiency, and cost-effectiveness in AI/HPC networking.

Why DriveNets Network Cloud-AI?

Performance

DriveNets enables the best Job Completion Time (JCT) performance at the highest scale, offering predictable, lossless back end cluster connectivity. With a proven 10%-30% improvement in JCT compared to any Ethernet-based solution, it ensures unmatched efficiency and reliability for demanding workloads.

Scalability

DriveNets provides a scalable, high performance solution designed to grow seamlessly with your business for up to 32K GPUs @ 800Gbps in a single cluster, delivering robust support for the ever-evolving demands of AI workloads.

Field Proven

The first and only scheduled Ethernet-based solution actively deployed and trusted by enterprise customers and hyperscalers worldwide.

Open and standard based

DriveNets solution is Ethernet-based, designed to seamlessly support an open ecosystem of partners. It offers unparalleled flexibility by accommodating any white box router, any optics, any NIC, and any GPU, enabling diverse deployment options. This approach ensures reduced total cost of ownership (TCO) while maintaining scalability and performance.

Simplicity

Simplicity lies at the heart of our solution, offering a true plug-n-play experience with no need for fine-tuning during setup, unlike traditional InfiniBand (IB) systems. Whether it's during bring-up or adapting to workload changes, the process remains seamless and efficient. Built on Ethernet, it eliminates the need for unique or specialized skills, making it accessible for teams without requiring additional knowledge or complex training.

Customer Production Cluster Results

Network Cloud-AI demonstrated improved JCT performance over standard Ethernet in all types of collective communications (performance improvement varies according to collective communication type).

JCT Improvement Per Collective Communication

1,280 xPUs + DriveNets Network Cloud-AI

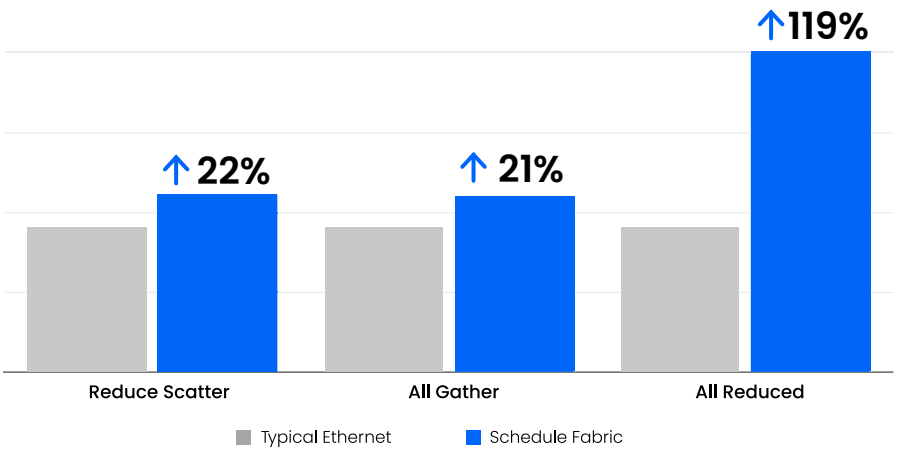

Source : ByteDance presentation in OCP October 2024
"Insights from Production: Scheduled Ethernet Fabric in Large AI Training Clusters"