DRIVENETS

AI-Grade Performance with DriveNets Network Cloud-AI

The roots of large-scale computing clusters for complex problem-solving can be traced back to the mid-20th century. However, the adoption of high-performance computing (HPC) for AI-driven workloads has surged in the past decade. Hyperscalers and HPC cluster builders have traditionally relied on solutions, utilizing their GPUs and InfiniBand-based connectivity for their high-performance capabilities, leading to excellent GPU utilization and job completion time (JCT) results.

Yet AI cluster builders want to move away from NVIDIA's proprietary solutions due to vendor lock and limited flexibility. They advocate for open standards like Ethernet and broader vendor support. However, to shift from InfiniBand, a new solution must deliver three key characteristics:

- **AI-grade performance with best JCT results**
- **Proven success in real-life, large-scale deployments**
- **Support of standard Ethernet**

AI-Grade Performance

- Best JCT results
- Predictable, lossless, low tail-latency connectivity, ensuring the highest GPU utilization

Field-Proven Solution

- Deployed in production by leading hyperscalers, NeoClouds, and enterprises
- Powering AT&T's core IP network since 2017

Standard Ethernet Support

Ensuring interoperability and innovation

DriveNets AI/HPC Fabric-Scheduled Ethernet (FSE)

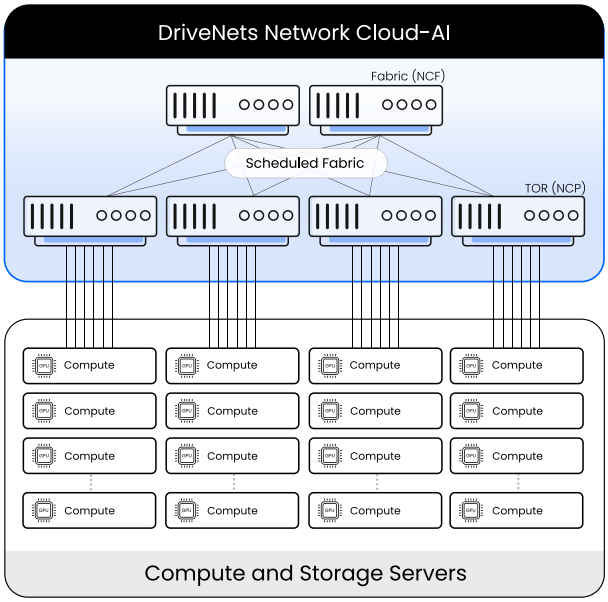

The DriveNets Network Cloud-AI solution delivers industry-leading AI networking fabric performance at scale, supporting up to 32,000 GPUs in a single cluster. Built on DriveNets' Fabric-Scheduled Ethernet (FSE) architecture, the solution ensures predictable, lossless backend connectivity, delivering significant improvement in JCT for AI workloads.

DriveNets' FSE is built on Distributed Disaggregated Chassis (DDC) architecture, an emerging standard defined by the Open Compute Project (OCP). This architecture uses standard white box switches in a Clos topology, efficiently scheduling and managing data flow to optimize network paths, eliminate congestion, and minimize latency and jitter. The result is a more predictable and lossless network connection.

3 Key Capabilities Driving AI-Grade Performance

1. **Cell spray:** Segmentizes data packets into uniform cells, distributed across all fabric device ports. Cells are reassembled at the egress leaf, ensuring optimal load balancing and the maximizing of fabric utilization.
2. **End-to-end virtual output queues (VOQs):** Mitigates congestion by queuing data packets at the source and transmitting them only when the destination is ready. This prevents head-of-line blocking and packet drops - common in traditional architectures.
3. **Zero-impact failover:** Using hardware indicators at the leaf level, DriveNets Network Cloud-AI detects physical failures in real time. Calculating each failure impact at the leaf level, this enables microsecond-level recovery with minimal disruption to the data plane. The failure signal is actively propagated throughout the cluster, ensuring the entire AI-scheduled fabric system dynamically adapts to prevent packet loss.

Cell spray, end-to-end VoQs and zero-impact failover work together to deliver the industry's best load balancing, congestion management, and high availability capabilities. All this is done without the need for fine tuning or complex configuration.

Network Cloud-AI additional benefits

Multi-Tenancy
- Isolated traffic flows eliminate noisy neighbor issues
- No overlay penalty
- Empowers any model, built to handle dynamic workloads

Highest Performance
- Best GPU utilization
- Predictable & lossless connectivity
- Proven high performance in enterprise, NeoCloud & hyperscalers

Openness
- Avoids vendor lock – Any GPU, Optics, NIC
- UEC ready
- Easy interoperability & innovation with standard Ethernet

Scale Across Multi-Sites
- Lossless multi-site connectivity
- Up to 80 km between sites

Fastest Deployment
- Fast & easy deployment with automated orchestration and ZTP
- No manual tuning required – best performance out-of-the-box

Performance Results in Customer Production Cluster

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

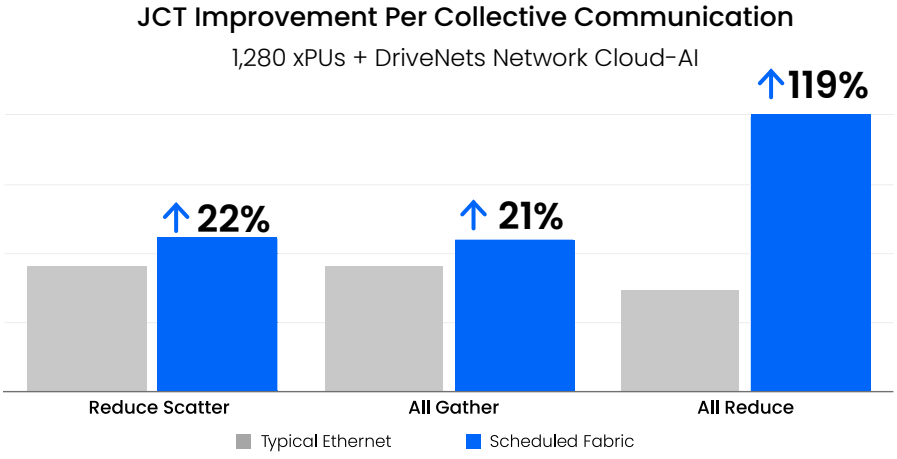

Source: ByteDance presentation in OCP (October 2024): "Insights from Production: Scheduled Ethernet Fabric in Large AI Training Clusters"

Learn more about how the DriveNets AI/HPC networking solution can help to offer more from AI workloads.