

DriveNets Network Cloud

DriveNets is a rapidly growing software company that has created a radical new way to build high scale networking infrastructures.

Most service provider networks and AI networking infrastructures are hardware-centric, based on monolithic systems, and, in some cases, on proprietary technologies, such as InfiniBand.

DriveNets Network Cloud and DriveNets Network Cloud-AI are software-based networking solutions based on a cloud-native architecture. Combining the scalability of standard Ethernet Clos with the high performance of a router chassis, DriveNets Network Cloud-AI delivers the highest networking performance and scale at the most optimal cost structure.

DriveNets Network Cloud software runs over standard white box hardware supplied by multiple ODMs, elastically scaling network capacity by adding additional white boxes into physical network clusters. Its software architecture ensures that all these white boxes act as a single elastic lossless network fabric.

DriveNets Network Cloud-AI

Large AI systems support large workloads running over individual compute and storage nodes that work together as a logical cluster and are connected via a high-capacity interconnect fabric, also known as AI networking. These networking solutions need to evolve to maximize the utilization of costly AI compute resources and support standard connectivity that enables vendor interoperability.

DriveNets Network Cloud-AI offers the highest performance AI networking at scale – up to 32K GPUs (800 Gbps) per cluster. Based on the largest-scale DDC (Distributed Disaggregated Chassis) architectures in the world, Network Cloud-AI provides predictable, lossless back-end cluster connectivity and 10%-30% proven improvement in JCT (job completion time) of high-scale, high-performance AI workloads.

DriveNets Network Cloud-AI maintains a GPU-, ASIC- and ODM- agnostic architecture and supports over 100% return on investment (ROI) at day one. It is now available for Broadcom's Jericho 2C+ and Jericho 3-AI based white boxes, supporting the highest performance Ethernet-based AI networking solutions today.

DriveNets at-a-glance

- **Founded:** 2016
- **Headquarters:** Israel
- **Vision:** Accelerate networking innovation to empower a future of new connected experiences
- **Founded: Total to date** \$587 million

Trusted Solution

- Based on Distributed Disaggregated Chassis (DDC) by OCP
- Running the world's largest network – AT&T's IP core
- Standard Ethernet
- Up to 32K GPUs (800G) per cluster
- 10%-30% JCT improvement

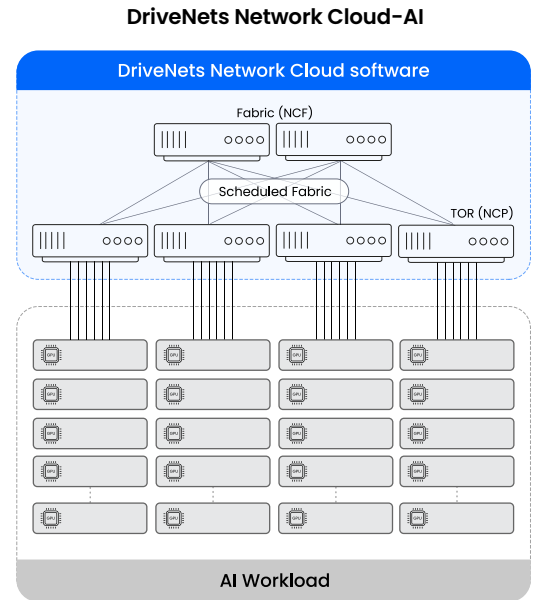


"ByteDance has been partnering with Broadcom and DriveNets to test the Scheduled Fabric over the past few months. The results have been promising, motivating ByteDance to deploy the world's first 1K xPU production cluster powered by Scheduled Fabric in July 2024"

www.broadcom.com/blog/peek-into-the-first-scheduled-ai-fabric-in-production

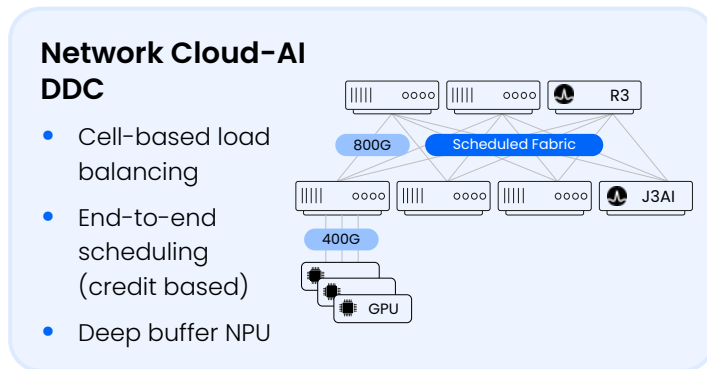
Highest performance Ethernet AI Networking

- **Highest JCT performance at highest scale:** predictable, lossless back-end cluster connectivity for up to 32,000 GPUs (800 Gbps) with proven 10%-30% JCT improvement vs. any Ethernet solution
- **Field-Proven:** Field-proven solution by top hyperscalers and an independent lab – Scala Computing (see results below)
- **Open & Flexible Tier-1 ecosystem:** Large industry ecosystem of leading hardware vendors (chip manufacturers, white box ODMs, and optical equipment producers)
- **Ultra-Fast Recovery:** Seamless failure mitigation, under 1ms using a hardware assisted detection mechanism



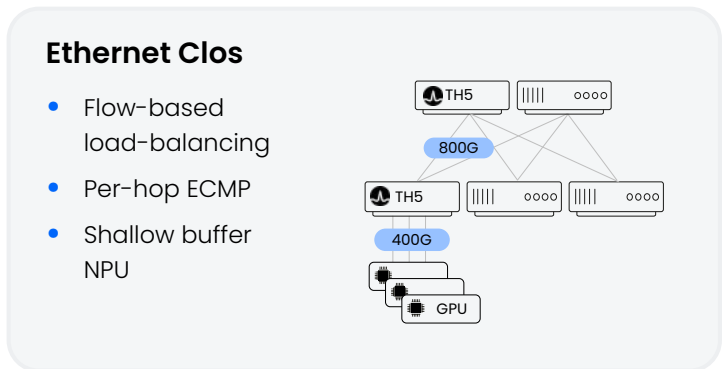
Validated by an independent lab – Scala Computing

Testing of Network Cloud-AI with 2K GPUs at 400G with multiple AI workloads demonstrates improved JCT performance by more than 10%.



Network Cloud-AI DDC


- Cell-based load balancing
- End-to-end scheduling (credit based)
- Deep buffer NPU



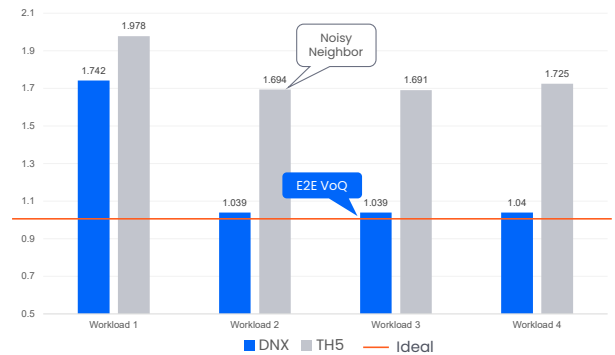
Ethernet Clos

- Flow-based load-balancing
- Per-hop ECMP
- Shallow buffer NPU

10% lower JCT for the DDC platform



50% NIC impairment affects all workloads in the Ethernet Clos



Network Cloud-AI is also optimized for GPU noisy neighbor scenarios, ensuring no performance impact on other AI jobs running in parallel. This is not the case with Ethernet Clos solutions where all AI jobs running on the same node are impacted. A GPU noisy neighbor scenario occurs when the performance of one or more GPUs is negatively affected by the activity of other GPUs on the same node. This can happen for a variety of reasons, such as resource contention on the network or performance degradation on the NIC.