PARLON NETWORK PATH SOLUTION

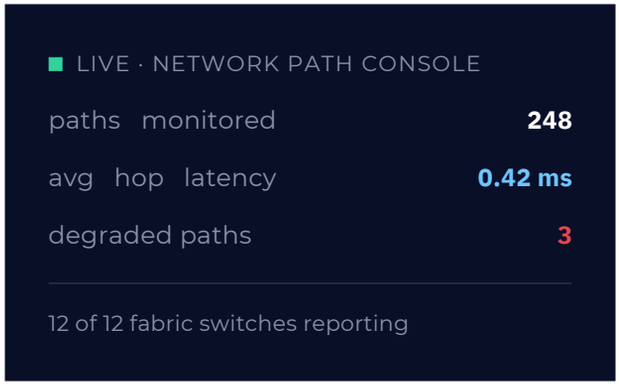
The path between your GPUs is invisible until a training job slows down and nobody knows why.

The Parlon Network Path Solution gives infrastructure and platform teams a live, hop-by-hop view of east-west traffic and fabric health across the AI cluster: GPU to GPU, not just north-south to the internet.

THE PROBLEM

Traditional network monitoring watches north-south traffic: the path from a user to an application. AI workloads live somewhere else entirely. A training job's real bottleneck is often east-west, GPU to GPU, node to node, across a fabric most monitoring tools were never built to see hop by hop.

A single degraded link, a routing change, a PFC storm on one switch, can quietly slow a distributed training job or an inference cluster long before anything alerts. By the time utilization drops enough to notice, the job has already been running slower for hours.



Why this matters for the AI buildout: a degraded fabric doesn't show up as a network alarm first. It shows up as a GPU utilization problem nobody can explain from the GPU dashboard alone. Reported GPU utilization already runs 47–68% in production fleets, and a fabric issue is one of the least visible reasons why. Seeing the path is what makes a GPU fleet's utilization numbers explainable, not just visible.

WHAT IT DOES

What the Parlon Network Path Solution does

HOP

Hop-by-hop path visibility

Every path between GPUs and nodes mapped hop by hop, not summarized into a single east-west average that hides where the slowdown actually is.

PFC

Fabric health signals

PFC storms, ECN marking, and lossless drops surfaced as they happen, correlated to the paths and jobs they affect.

FLOW

East-west NetFlow

Traffic between GPUs and nodes analyzed the same way north-south traffic always has been: top talkers, protocol mix, and traffic heatmaps.

ALRT

Path divergence alerts

Routing changes that quietly degrade a path get flagged before they cost a job hours of throughput.

35%

WHY IT MATTERS

of teams believe their tools are ready to observe AI-era failure modes, including the east-west traffic moving between GPUs and nodes.

Source: EMA, 2026

DEPLOYMENT & COVERAGE

One fabric, one data model.

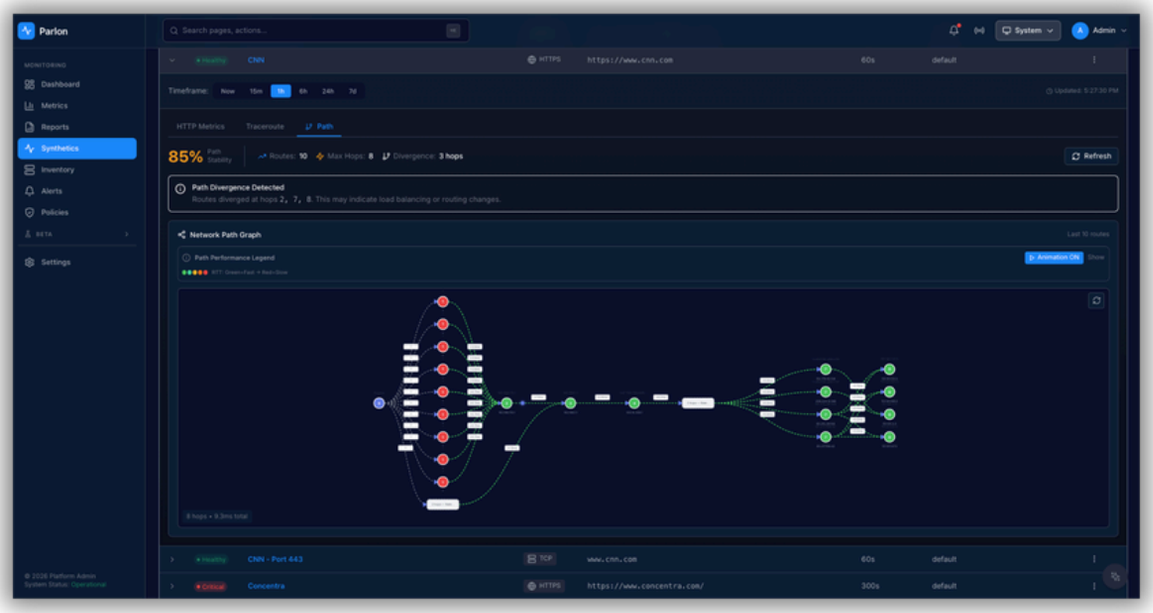
Network Path telemetry flows through the same Normalization Engine as every other infrastructure source, and deploys wherever Parlon does: SaaS, on-premises, hybrid, or fully air-gapped.

NetFlow/IPFIX

P2P: ICMP, Traceroute, Throughput

PRODUCT

Every hop, every path, on one screen



The Network Path console: hop-by-hop path map, fabric health signals, and east-west traffic analysis in one view.

LEARN MORE