CORPORATE SOLUTION GUIDE

Observability for AI infrastructure.

GPU fleets, power, agent activity, and the network between them — seen and governed in one platform. The infrastructure layer every AI workload runs on, and the layer legacy tools were never built to see.

THE SHIFT

AI workloads run on physical and operational layers most observability platforms don't reach: the GPU fleet, the power feeding it, the agents now operating inside it, and the network moving data between all three. Most platforms approach AI observability from the software side and stop at the metric. Parlon starts at the silicon and the fabric.

That's not a positioning choice. It's where the buildout is actually running into trouble first — and where nobody has built the tooling yet.

47.7%
of enterprises already run AI training or inference workloads today
(EMA, 2026)

35%
believe their tools are ready to observe AI-era failure modes
(EMA, 2026)

73%
of IT pros likely to replace an observability tool within two years
(EMA, 2026)

THE PRODUCT

Four solutions, one platform

GPUs, power, agents, and the network between them — the four atomic units of the AI infrastructure layer, each shipping today, each backed by its own solution guide, all running on one normalized data model.

GPU FLEETS
Utilization, thermal throttle, and ECC/XID errors, down to a single device.

POWER & PDU
Breaker-level load, A/B feed redundancy, and capacity runway before a rack trips.

MCP & AGENT ACTIVITY
Every AI agent read, attributed and audit-logged.

NETWORK PATH
East-west traffic and fabric health across the AI cluster, hop by hop.

Legacy platforms watch applications. Almost none treat GPU fabric, power, and agent governance as first-class products. This is the ground the giants don't own — and where Parlon's depth already sits.

THE FOUNDATION

Why this is credible, not just another dashboard

Every wedge above runs on the same foundation Parlon has already built and deployed: the legacy-replacement engine that makes AI infrastructure observability possible.

NORM

Normalization at ingest

Every source mapped to a unified schema the moment it arrives. Vendor detail preserved. Correlation is instant, not an afterthought.

SYN

Synthetics + telemetry, unified

Active synthetic testing and continuous telemetry in one native system, including LLM-aware workflow checks. No modules, no swivel chair.

TUNE

Alert Auto-Tune™

A human-approved workflow that recommends thresholds with evidence, reducing alert noise. When Parlon alerts, it's worth acting on.

DEPLOY

Deploy anywhere

SaaS, on-premises, hybrid, and air-gapped deployments with customer-controlled boundaries including fully sovereign environments.

Parlon is observability built for AI infrastructure.

It sees the GPUs, fabric, and power under your AI workloads, governs the agents watching them, and runs self-hosted with predictable economics. Observability funds the business today. Governed AI operations compounds it.

Let's talk about the AI infrastructure layer.

LEARN MORE