

AI Factory Observability for High-Frequency Trading

Turn AI infrastructure into trading advantage.

Every market event creates a race. Market data arrives, AI models generate signals, execution algorithms make decisions, and orders reach exchanges in microseconds. Success depends on an execution system spanning AI models, GPU clusters, FPGA infrastructure, market-data pipelines, ultra-low-latency networks, storage, and applications working together without introducing delay or variability.

Legacy observability stops at individual components. GPU monitoring stops at utilization. APM stops at code. Infrastructure monitoring stops at hardware. None understand the complete AI-powered trading system or identify the system constraint driving degraded execution performance.

The Virtana AI Observability Platform delivers AI Factory Observability by continuously discovering, correlating, and modeling every dependency across the trading system. The result is a live operational model that enables teams and system-aware AI agents to identify the system constraints affecting trading performance before they impact the market. Rather than generating more alerts, Virtana correlates every dependency—from AI inference to exchange connectivity—into a single operational context that enables teams and system-aware AI agents to identify the actual system constraint affecting trading performance.

One Trading System. One Operational Model.

High-frequency trading is no longer an application running on infrastructure. It is a distributed AI execution system.

Virtana continuously discovers and correlates relationships across AI workloads, GPUs, FPGA infrastructure, market-data pipelines, storage, networking, Kubernetes, and execution services into a live System Dependency Graph. Every trace, metric, log, event, and hardware signals are unified into one system, giving AI complete operational context for high-velocity trading.

This operational model becomes the foundation for everything Virtana does. Every investigation, recommendation, and AI-generated insight is grounded in this live operational model, eliminating manual correlation across disconnected tools. Instead of asking teams to manually correlate dashboards, logs, metrics, and traces, Virtana understands how every dependency contributes to execution quality, infrastructure efficiency, and AI workload performance.

System-Aware Agentic AI That Reasons Across the Entire Trading System

The Virtana AI Observability Platform doesn't simply surface alerts—it equips system-aware AI agents with complete system context. Because every dependency is continuously modeled in the System Dependency Graph, Virtana's agentic AI investigates incidents in the moment they occur, traversing relationships across AI models, GPU infrastructure, applications, storage, networking, market-data pipelines, and execution services to identify the actual system constraint.

Rather than asking operators to assemble war rooms, Virtana delivers evidence-backed answers that explain:

- What changed
- Where the constraint originated
- Which trading services are affected
- The downstream impact on execution quality, latency, fill rates, and AI infrastructure
- Recommended remediation actions

Optimize the AI Trading Factory

For modern trading firms, AI infrastructure is no longer simply a technology investment—it's part of the trading strategy.

Virtana helps firms maximize AI infrastructure investments by:

- Increasing GPU utilization and infrastructure efficiency
- Identifying hidden system constraints before they impact execution quality
- Optimizing AI economics through better capacity planning and resource utilization

Rather than monitoring isolated infrastructure, Virtana helps organizations optimize the entire AI trading factory.

Operate Through Natural Language and Agentic AI

Virtana makes the operational model available to both people and AI.

Virtana enables quantitative developers, SREs, platform engineers, and trading operations teams to investigate the environment using natural language.

Teams can simply ask:

- Why did inference latency increase after market open?
- Which dependency reduced fill rates?
- What infrastructure constrained today's trading models?
- Why did execution latency increase during peak market volume?

Virtana delivers evidence-backed answers grounded in live system context—not statistical guesses. The Virtana Model Context Protocol (MCP) Server extends this same operational model to enterprise AI agents and automation platforms, enabling organizations to integrate trusted system context into agentic workflows without sacrificing governance or explainability.

Why Virtana

High-frequency trading is no longer an application running on infrastructure.

It is an AI-powered execution system.

- Legacy APM stops at code.
- Infrastructure monitoring stops at hardware.
- GPU monitoring stops at utilization.

Only the Virtana AI Observability Platform continuously understands the complete execution system behind every AI-powered trading decision.

Virtana transforms fragmented telemetry into system understanding. By correlating every dependency into a live operational model and applying system-aware AI, the platform identifies the constraints limiting execution quality, AI infrastructure efficiency, and trading performance.

Legacy monitoring shows fragmented symptoms. Virtana identifies the system constraint—and enables teams to act with confidence before trading performance is affected.

About Virtana

Virtana delivers the deepest and broadest observability platform for hybrid and multi-cloud environments, with full-stack AI observability spanning applications, services, data pipelines, GPUs, CPUs, networks, and storage. Powered by high-fidelity data and agentic AI, Virtana provides unmatched visibility across end-to-end IT services and AI workloads, correlating health, performance, cost, and user impact in real time. With advanced event intelligence and autonomous insight generation, Virtana delivers clarity no other provider can match. Trusted by Global 2000 enterprises and public sector organizations, Virtana helps IT operations and DevOps teams reduce risk, strengthen resilience, improve efficiency, and modernize with confidence across multi-cloud, on-premises, and edge environments.

Learn more at virtana.com.

info@virtana.com | **+1 408-579-4000** | **virtana.com**

©2026 Virtana. All rights reserved. Virtana is a trademark or registered trademark in the United States and/or in other countries. All other trademarks and trade names are the property of their respective holders.