



NVIDIA Minerva

The AI-native AIOps platform
for maximizing AI factory yield.



Transforming AI Factory Operations

AI factories are among the most complex systems ever built, requiring operators to manage hundreds of thousands of GPUs, high-performance networking fabrics, schedulers, and distributed AI workloads around the clock.

At this scale, traditional data center operating models break down. Teams face continuous alerts, scarce expertise, and the challenge of identifying the root cause of performance issues across millions of interacting components. Problems are often subtle: a GPU running hotter than its peers, an NVMe drive degrading under load, or a switch generating pause frames. Individually, these anomalies appear insignificant; collectively, they create stragglers, failed jobs, unpredictable latency, and lost tokens.

The result is slower training, lower token production, and higher operating costs. Maximizing AI factory yield requires an end-to-end platform that automates operations, correlates signals across the full stack, and keeps infrastructure performing at peak efficiency.

The Solution: NVIDIA Minerva

NVIDIA Minerva™ is an intelligent, AI-native, end-to-end AIOps software platform purpose-built to maximize the availability and performance of AI factory infrastructure. As part of the full-stack NVIDIA Vera Rubin Platform, NVIDIA Minerva harnesses metric, log, and topology telemetry from both infrastructure and AI workloads and applies a broad range of purpose-built ML algorithms and agentic techniques. NVIDIA Minerva turns a flood of raw telemetry into proactive, automated action that maximizes AI factory yield, including GPU hours, goodput, cluster uptime, and token production.

Comprehensive, Full-Stack Visibility

NVIDIA Minerva's strength begins with what it can see. It correlates telemetry across GPUs, NVIDIA Spectrum-X™ Ethernet (for scale-out, scale-across, and scale-in), NVIDIA NVLink™ (for scale-up), SuperNICs, DPUs, NCCL, schedulers, and AI workloads. The toolset transforms this information into a dynamic knowledge graph that maps the cluster and the jobs running on it.

Key Features

- > AI Factory End-to-End Health
- > Job-Centric Observability
- > Supervised Log Errors
- > Zero-Day Log Anomaly Detection
- > ML-Based Anomaly Identification
- > Auto-Correlation and Noise Reduction
- > Automated Predictive Maintenance
- > Built-in AI Assistant
- > MCP Server Connection to External AI Tools
- > Out-of-the-Box Dashboards and Alerts

Key Benefits

- > Maximize Effective GPU Hours
- > Unified View of Infrastructure and Workloads
- > Reduce Operational Noise
- > Move From Reaction to Prevention
- > Increase Mean Time Between Interrupts (MTBI)
- > Decrease Mean Time to Repair (MTTR)
- > Flexible Deployment Models




Figure 1: NVIDIA Minerva software breakout.

This full-stack context enables NVIDIA Minerva to connect a single anomalous interface among hundreds of thousands to the specific job it affects. As the primary network observability solution for NVIDIA AI factories, NVIDIA Minerva bridges infrastructure and workloads in ways that generic observability tools and siloed point solutions cannot.

NVIDIA Minerva analyzes billions of data points per hour in real time using models purpose-built for AI factory infrastructure and workloads. It learns normal behavior, predicts anomalies before they become incidents, and auto-correlates events across layers, enabling operators to focus on fewer, more relevant issues faster. Built into the platform are Sherlock, an AI-native agentic assistant, and an MCP server that enables operators to triage and debug issues through natural-language conversations grounded in live telemetry. Together, they democratize scarce expertise across operators of all experience levels.




Figure 2: Sherlock displaying network errors via natural language prompts.

The NVIDIA Minerva stack includes the NVIDIA Minerva Collector and the NVIDIA Minerva Platform. The NVIDIA Minerva Collector gathers telemetry across multiple standards, including OpenTelemetry (OTEL), gRPC/gNMI, syslog, DCGM, and Slurm. The NVIDIA Minerva Platform includes the topology graph, known and unknown anomalies, ML autocorrelation, automated workflows, and an agentic framework that turns data into actionable intelligence.

Use Cases

Proactive Maintenance

The dominant loss mechanism in AI factories comes from “gray failure”: components that are degrading but have not yet failed. NVIDIA Minerva catches these before they impact AI factory operations. It auto-discovers a graph model of the AI cluster across Spectrum-X Ethernet and NVLink, with hierarchical zoom that holds up at scale. It detects anomalies across GPUs, switches, SuperNICs, and DPUs, matching against known hardware, system, and SDK errors curated by NVIDIA experts and customers.

Time-machine filtering compresses investigation time from hours to minutes. Billions of telemetry data points become high-fidelity, actionable alerts about critical events, and scheduler integration turns those alerts into automated mitigation, including proactive node draining and maintenance.

The result is advanced warning on emerging failures before they reach workloads. Shifting from reactive troubleshooting to proactive maintenance reduces mean time to repair, prevents hundreds of lost GPU hours, and keeps AI workloads running at peak performance.




Figure 3: NVIDIA Minerva dashboard displaying AI cluster health.

Agentic Debugging

NVIDIA Minerva enables agentic, topology-wide root cause analysis for AI factories. Its native AI debugging assistant, Sherlock, is backed by live telemetry and the cluster knowledge graph. Using a natural-language interface, operators can query the state and health of AI factory infrastructure, debug complex issues, and receive recommended remediation. NVIDIA Minerva’s MCP server exposes the same capabilities to external agentic tools, including Claude and OpenClaw.

These capabilities fit into existing operational workflows and democratize scarce AI factory expertise. With agentic debugging, a Tier 1 engineer can reach a conclusion that previously required an NVIDIA specialist.




Figure 4: NVIDIA Minerva MCP server for integration with Claude code.

Full-Stack Correlation

AI factories span an enormous range of interconnected systems. When an issue occurs, its root cause may originate in one layer and surface as a symptom in another.

Generic observability platforms monitor either infrastructure or applications. NVIDIA Minerva spans both, revealing the relationships that matter: how a network anomaly affects GPU utilization or how a scheduler event contributes to degraded job throughput.

NVIDIA Minerva filters telemetry by job ID for workload-specific visibility, correlating individual jobs with the logs, metrics, errors, and anomalies beneath them. A single dashboard surfaces critical errors across both metrics and logs, and nodes with critical errors can be drained directly from the job scheduler. Operations teams get one unified view of the AI factory instead of fragmented pieces from siloed tools that make cross-layer debugging slow and costly.

Noise Reduction

Alert fatigue is one of the largest hidden costs in AI factory operations. A network issue that cascades across hundreds of components can generate thousands of alerts. NVIDIA Minerva's topology-aware correlation engine understands which components depend on which, so a cascading fault is presented as one incident with a root cause rather than hundreds of symptoms. NVIDIA Minerva then delivers that single prioritized signal, with actionable context, through tools such as Slack or ServiceNow. The result is a significant reduction in notifications and much faster resolution time. Teams spend less time triaging alerts, diagnose issues faster, and focus on the highest-impact operational problems.




Figure 5: Minerva detects an anomaly.

Network Observability

As the primary network observability solution for NVIDIA AI factories, NVIDIA Minerva includes out-of-the-box dashboards and alerts built for NVIDIA Spectrum-X Ethernet scale-out, scale-across, and scale-in; NVLink scale-up; hardware health; and network services software. Dashboards align with NVIDIA Reference Architectures, giving operations teams best-practice visibility from day one.




Figure 6: NVIDIA Minerva displaying a network topology.

Adding AIOps Capabilities to DSX OS

NVIDIA DSX™ is an AI-factory-scale platform that unifies design, simulation, and operations across the full stack through a common co-designed architecture. Within the DSX platform, DSX OS is the operating layer for running and scaling multi-tenant AI factories.

NVIDIA Minerva Collector is part of DSX OS and provides customers with an NVIDIA Reference Architecture-compliant telemetry collection from network infrastructure. Integration with DSX OS lets operators correlate infrastructure telemetry with the broader DSX operational surface, including facility-level signals exposed by other DSX OS components, for a single coherent view from GPU to grid.

DSX OS customers can add NVIDIA Minerva as an option to bring its AI/ML insights into the DSX OS observability stack. NVIDIA Minerva can deliver value inside a DSX OS deployment and as a standalone platform for teams not yet running DSX. Either way, customers get the operational intelligence that maximizes AI factory yield, plus a clear path to the rest of the DSX ecosystem as their environment grows.

Consumption Models

NVIDIA Minerva provides multiple consumption models to fit into an AI factory's existing workflow. These models range from basic DIY observability using the NVIDIA Minerva Collector, to using the Collector with the NVIDIA Minerva Platform's observability and anomaly detection features alongside external observability solutions, to full turnkey deployments.

Platform Specifications

Technical Specifications

Telemetry Sources	GPUs; NVIDIA Spectrum-X scale-out, scale-across, and scale-in; NVLink scale-up switches; SuperNICs; DPUs; hosts/nodes; storage NCCL; and schedulers.
Ingestion Standards	OpenTelemetry (OTEL), gRPC/gNMI, syslog, DCGM, and Slurm.
Analytics and ML	Purpose-built supervised and unsupervised ML models, NLP-based log anomaly detection, prediction models, and a proprietary topology-aware auto-correlation engine.
Knowledge Graph	Auto-discovered, dynamic graph model of infrastructure and jobs spanning single or multiple data centers, enriching telemetry with real-time context.
Agentic Interface	MCP server and platform-native multi-agent assistant; integrates with agentic tools and ecosystem AI agents.
Remediation and Actions	LLM-recommended remediation; human-in-the-loop or automated mitigation via scheduler, SuperNIC, and switch APIs and external automation systems, such as Ansible.
Integrations	Notifications to Slack, Microsoft Teams, and PagerDuty; ticketing via ServiceNow; visualization via Grafana; scheduler integration, including Slurm.
Enterprise Readiness	RBAC, SSO, multi-tenancy, high availability, and proven hyperscale operation across deployments with hundreds of thousands of GPUs.
Deployment	Fully on-premises; horizontally scalable microservices across a collector stack and a platform stack.

Ready to Get Started?

To learn how NVIDIA Minerva can maximize the yield of your AI factory, contact your NVIDIA representative:
nvidia.com/networking-contact-sales

©2026 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, NVIDIA Minerva, NVLink, and Spectrum-X are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective owners with which they are associated.5730117. OCT26

