


# Full-Stack Observability for Full-Speed AI

Unify visibility, control, and resilience across every layer of your AI infrastructure

## See more and operate faster

Explore how CoreWeave Mission Control™ unlocks greater speed, resilience, and reliability for your most complex AI workloads.

[Learn more](#)



## A new operating standard for observable AI infrastructure

Access metal-to-model performance metrics, verifiable control, and real-time audit telemetry all from one unified system for comprehensive AI-native observability.

### Most clouds only offer near-sighted solutions

General-purpose monitoring tools weren't built for the speed, complexity, or scale of modern AI.

- ✗ **Fragmented visibility** across infrastructure and workloads
- ✗ **Slow root cause analysis** for distributed training, often requiring manual intervention
- ✗ **Lack of verifiable auditability** and compliance confidence that puts security at risk

These systems collapse under pressure, leading to wasted GPU hours, stalled training runs, and unpredictable costs.

### CoreWeave offers a purpose-built foundation

Built for today's demanding workloads, CoreWeave introduces a new operating standard for observability into workloads.

- ✓ **Greater transparency** into what's happening across your clusters as it happens
- ✓ **Enhanced reliability** that ties observability and health checks to fast triage
- ✓ **Deeper insights** into GPU- and rank-level performance analytics

This is observability designed for the next generation of AI infrastructure: clearer, faster, and ready for production at scale.

### COREWEAVE CLOUD GIVES YOU

**+100k**

metrics tracked per run

**1 million**

data points ingested per second

**↑96%**

goodput

# Unlock resiliency with end-to-end visibility

CoreWeave provides vertically integrated observability from the application layer down to bare metal, which helps run AI workloads with greater speed, resilience, and reliability. Get a full view of detailed metrics and data available for your clusters. Out-of-the-box. Zero setup. Zero extra charges.

## A PURPOSE-BUILT FOUNDATION



### 01

#### Comprehensive dashboards

Monitor your entire fleet of NVIDIA GPUs in one place, complete with details about every node, via fully managed, curated Grafana dashboards. Plus, surface telemetry and remediation guidance through conversational workflows with CoreWeave Mission Control™ Agent.

### 02

#### Fast access to granular metrics

Unlock deep bottleneck analysis with GPU Straggler Detection. Access near-real-time node, network, and storage metrics—including temperature, power consumption, ingress and egress traffic, and more— to ensure your cluster is fully and efficiently utilized.

### 03

#### Full compliance and auditability

Maintain audit-ready visibility through a fully managed Loki API that makes logs immediately available, eliminating the need for rehydration from cold storage.

### 04

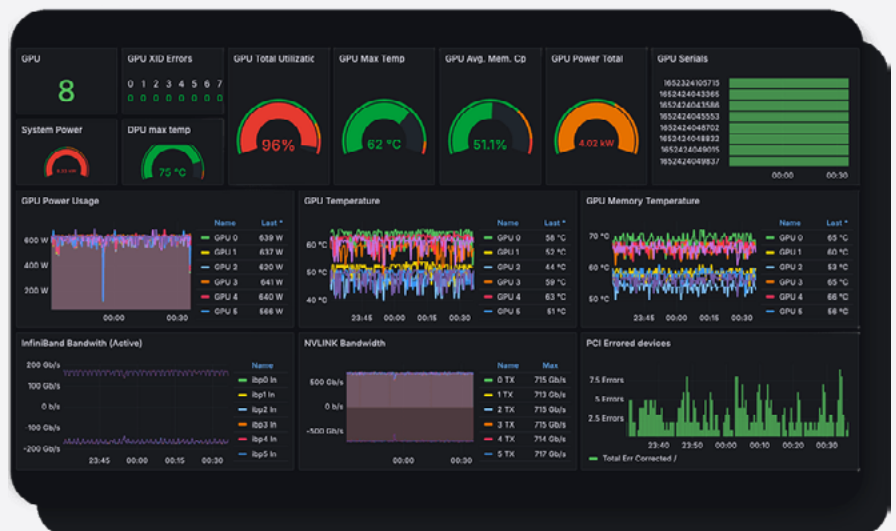
#### No-cost telemetry forwarding

Experience fully managed telemetry forwarding with CoreWeave Telemetry Relay, delivering encrypted, reliable, and filtered log streams directly to your external data platforms or on-prem endpoints.

FIGURE 1

#### You see what we see

Out-of-the-box, curated Grafana dashboards intuitively reveal what's going on in your clusters in near-real-time, so you can keep training jobs on track and running at full efficiency.



POWERING AI INNOVATION



## A unified system for AI reliability, transparency and insight

CoreWeave Mission Control™ integrates observability data, secure audit visibility, and expert-led automated operations into a single, cohesive operating system. This is reliability and transparency that unlock the full potential of your AI breakthroughs.

### CoreWeave Observability


- ✓ 100K+ metrics tracked per run
- ✓ 1M+ data points ingested per second via Weights & Biases integration
- ✓ Curated Grafana dashboards of massive-scale AI clusters
- ✓ Real-time ingress and egress traffic
- ✓ SUNK integration for Kubernetes and Slurm job telemetry

### CoreWeave Telemetry Relay

- ✓ Real-time audit and access telemetry
- ✓ Forward logs across all CoreWeave services
- ✓ Encrypted, buffered, strict SLOs
- ✓ Multi-destination routing

### GPU Straggler Detection

- ✓ Rank-level NCCL insights
- ✓ Slowdown detection “in minutes,” not hours/days
- ✓ GPU- and node-level diagnosis
- ✓ Alert recipes and Grafana overlays
- ✓ Automated fleet lifecycle controllers reducing MTTD/MTTR





## See more clearly. Build more confidently.

CoreWeave Mission Control™ establishes a new operating standard for AI, combining security, talent services, and observability—so you get the insights and reliability you need to move at pioneering pace.

### THE NEW OPERATING STANDARD



Up to 40% lower costs with measurably higher GPU utilization

As the Essential Cloud for AI, CoreWeave delivers a truly purpose-built stack that gives you end-to-end visibility into your AI workloads and the clusters they run on.

With deep observability and robust cluster health management, your team sees measurably higher GPU utilization and up to 40% lower costs.

## Move faster with deeper insights

Explore more benefits of a purpose-built AI cloud powering your most complex workloads.


Contact Us



### GPU Straggler Demo

Watch the CoreWeave Mission Control™ agent in action.

Watch →



### Automated User Provisioning

Accelerate cluster access for Slurm with AUP and SUP.

Read more →



### Explore CoreWeave Mission Control™

Reliability, transparency, and insight for large-scale AI workloads.

Learn more →