

CoreWeave Mission Control: The Operating Standard for the AI Cloud

Run AI at production scale with cloud-level reliability, transparency, and insights you can act on

Book a Mission Control review

Map your environment to the Mission Control operating standard and get a tailored rollout plan.

[Book a time](#)



Why AI demands an operating standard

AI scale has blown past what human operators and general-purpose clouds can reliably manage. When training or serving across hundreds or thousands of GPUs, the failure modes aren't theoretical—silent stragglers, audit gaps, and wasted compute become daily realities.

From fragmented signals and operations...

Distributed training and large inference systems depend on reliable coordination across GPUs, nodes, networks, and identity layers. But most clouds expose fragmented telemetry, slow remediation paths, and security signals scattered across services. Teams end up stitching together dashboards, exporters, and runbooks just to answer basic questions like “Why did this job stall?” or “What changed in the cluster?”

... to unified insights across the stack

Mission Control combines security, talent services, and observability in a single system. The tight integration gives teams real-time visibility, dependable fleet remediation, and audit-ready evidence of workload behavior. With capabilities like Telemetry Relay and GPU Straggler Detection, Mission Control delivers unmatched transparency and performance insight at scale—defining a new operating standard for AI cloud.

THE SCALE PROBLEM



Real-time visibility into GPU, network, and storage behavior



Continuous health checks and automated fleet remediation



Expert-backed and agentic operations tuned for AI-native workloads

Everything connected. Every insight actionable.

Mission Control is built from three integrated core capabilities: security, talent services, and observability—that define how AI workloads operate on CoreWeave Cloud.

01 Security

Mission Control brings CoreWeave's security foundation into a unified operational layer. IAM, RBAC, and Support Access Management give teams clear control over identity and access. Continuous audit logging provides traceability, and SIEM-ready telemetry delivery ensures security and compliance teams work from complete, real-time evidence.

02 Talent services

Mission Control's Talent Services combine automated lifecycle control with human expertise. Node and Fleet Lifecycle Controllers handle burn-in, qualification, and remediation to keep clusters healthy. When automation isn't enough, direct-to-expert support connects teams with the engineers who operate the platform.

03 Observability

Mission Control provides full-stack visibility, delivering metal-to-model telemetry across GPUs, nodes, network fabric, storage, Kubernetes, and APIs. Correlated dashboards surface job behavior and system signals in one place, making performance issues quickly diagnosable instead of time-consuming mysteries.


Together, these three core capabilities form the foundation of Mission Control—delivering the reliability, transparency, and insight pioneering AI teams need to operate at scale.

INTEGRATING TOOLS,
TEAMS, TECH



Up to 80% faster
model training speeds

FOR IBM



From architecture to outcomes

Mission Control translates system integration into operational speed—more stability, less wasted compute, and faster iteration at scale.

Continuous fleet monitoring enables fast issue resolution across clusters of any size, keeping long-running training on track. High qualification and lifecycle controls drive high goodput and sustained MFU, reducing wasted compute and stabilizing iteration speed. Correlated model and system telemetry—including integrations like Weights & Biases—dramatically shrinks diagnosis time. Mission Control plays a key role in supporting these outcomes by keeping fleets healthy and predictable. Teams get a complete signal-to-action loop: issues surface, meaning is clear, and remediation is built in.

96%

Training goodput

High effective compute driven by healthy fleets and fast recovery

3.66 days

MTTF

10x improvement in stability vs. 0.33-day industry baseline

50%

Fewer interruptions

Fewer interruptions per day thanks to proactive node lifecycle control and full-stack observability.

ENGINEERED FOR SCALE AND TRANSPARENCY

Up to 20% more FLOPs/ GPU/hr

FOR DATABRICKS

CoreWeave Mission Control

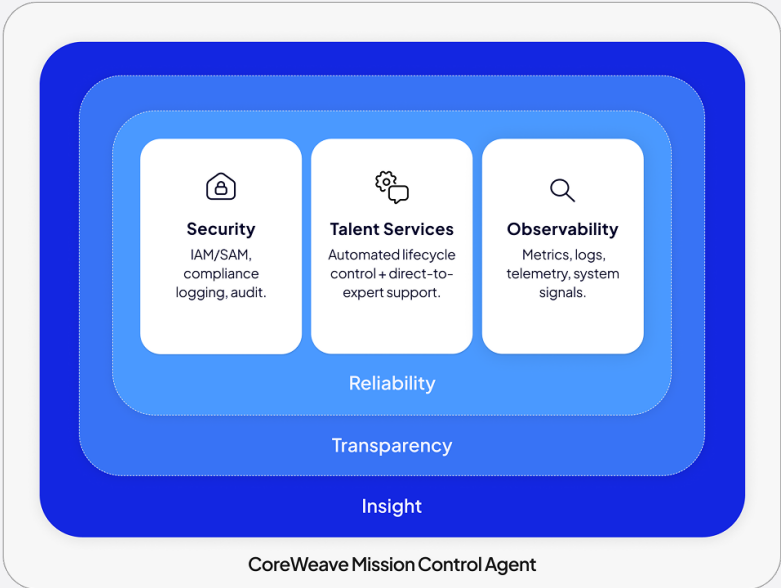


FIGURE 1

Mission Control runs through the entire CoreWeave stack. It unifies security, talent services, and observability into one consistent operating standard for AI workloads.

Telemetry Relay: SLO-backed telemetry delivery

Telemetry Relay is a fully managed path for delivering logs and metrics to external observability platforms with durability, buffering, and SLO-backed reliability. It eliminates custom pipelines and extends full-stack telemetry into existing SIEMs and monitoring tools.

CoreWeave provides deep, vertically integrated observability—GPU, node, network, storage, Kubernetes, and system signals—all available in CoreWeave Grafana and via PromQL and LogQL APIs. Telemetry Relay extends this by delivering selected logs and metrics to external platforms when needed, with durability, buffering, backpressure control, and strict completeness and latency SLOs.

Streams define selectable sets of telemetry. Forwarding endpoints specify where that data goes, including SIEMs, HTTP log ingestors, Prometheus remote-write targets, and more. Pipelines link streams to endpoints, making it easy to route audit logs or metrics to external systems without maintaining custom code or infrastructure.

Telemetry Relay gives security and platform teams SIEM-ready audit delivery with no custom scraping or exporters. It integrates seamlessly with existing toolchains and introduces a powerful new control-plane primitive: reliable, SLO-backed telemetry delivery at scale.

Teams can rely on Telemetry Relay to centralize audit visibility for compliance, unify SRE alerting with platform metrics, and integrate GPU, node, and fabric telemetry into existing observability stacks—all without operating separate pipelines.

TELEMETRY WITHOUT FRICTION

“

[Our models] were trained 100% on CoreWeave infrastructure. I think not being on that kind of infrastructure would’ve delayed us by at least a few months.”

Timothée Lacroix
CTO, Mistral AI

CoreWeave Telemetry Relay		Competitors	
✓	Fully managed observability platform	✗	Multiple services or DIY solution consolidating CSPs
✓	Telemetry forwarding	✗	Limited or no telemetry forwarding
✓	Audit and access log forwarding	✗	Limited audit and access log forwarding
✓	Robust SLOs	✗	Best-effort delivery only
✓	No additional costs for features	✗	Complex and unpredictable cost structure

GPU Straggler Detection: Rank-level insight for distributed training

GPU Straggler Detection exposes slow or degraded GPUs inside distributed training jobs—grey failures—using NCCL-level telemetry. It highlights the exact rank and node causing stalls and provides the surrounding signals teams need to diagnose issues quickly.

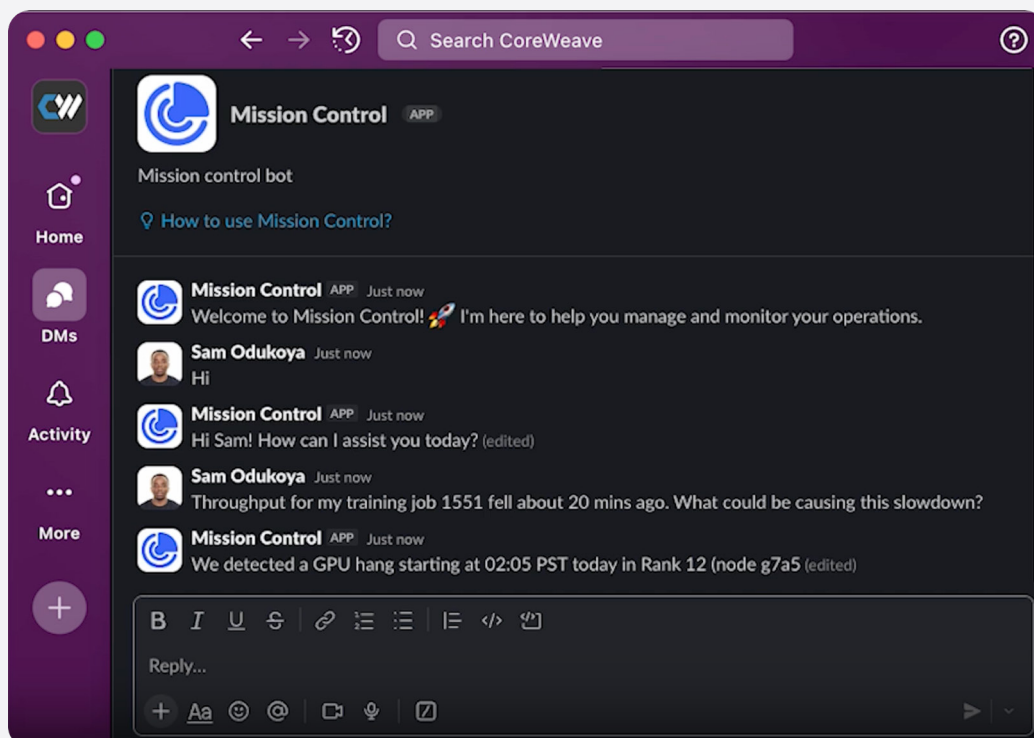
In distributed training, a single slow rank can stretch iteration times, waste compute, and stall progress, yet these issues are notoriously hard to pinpoint using traditional logs or high-level metrics.

GPU Straggler Detection analyzes NCCL-level telemetry to spot slow or divergent ranks during collective operations. It pinpoints the exact GPU and node responsible, and overlays rank-level bandwidth, latency, and message patterns with job state and system signals for precise root-cause analysis.

This eliminates one of the hardest debugging loops in distributed training. Teams no longer sift through logs or resubmit jobs repeatedly to isolate performance issues. Instead, they get clear, actionable insight that reduces wasted compute and accelerates iteration.

When Mission Control detects a straggler, it provides recommended actions, such as cordoning the node or rescheduling the job, to restore healthy performance quickly. This shrinks investigation time and keeps large-scale training runs on track.

DIAGNOSE STRAGGLERS AS THEY HAPPEN



Turn AI operations into measurable advantage

Mission Control translates unified operations into tangible results—the advantages leading AI teams rely on to move faster and waste less.

VALIDATED AT SCALE

Reliability: fleets that stay healthy

Continuous health checks, lifecycle control, and automated remediation keep GPU fleets productive and reduce infrastructure-related job disruptions.

Transparency: metal-to-model observability

CoreWeave Observe™ and Telemetry Relay deliver vertically integrated telemetry and SIEM-ready audit visibility so platform, security, and compliance teams operate from the same data.

Insight: from symptoms to fixes

GPU Straggler Detection and NCCL-level observability make slow or degraded ranks obvious, shrinking diagnosis time and reclaiming wasted compute.

Customer-proven at frontier scale

Leading AI labs rely on Mission Control to keep frontier-scale clusters stable, detect fabric and node issues early, and maintain predictable iteration schedules.


Future-ready foundation

Mission Control's architecture—durable telemetry queues, rich labeling, and expert-backed workflows—lays the groundwork for agentic operations, including the Mission Control Agent. The Agent interacts with Mission Control data and insights to explain incidents, answer questions, and propose actions.



Standardize how you run AI in production

AI workloads are only getting larger, more distributed, and more regulated. Standardize how you run them with CoreWeave Mission Control.

[Contact Us](#)


Form

Mission Control review

Map your environment to the Mission Control operating standard and get a tailored rollout plan.

[Book a time →](#)


Webinar

Experience AI-Native Observability

See how vertically integrated observability keeps AI workloads fast, stable, and efficient.

[Watch on demand →](#)


Digital

Keep exploring

Learn more about the operating standard for AI.

[Take the next step →](#)