Control Room Dashboard

Real-Time Command and Control for GPU-Accelerated Research Operations

RICCHE LTD

ricche.ai

Powered by the NVIDIA Accelerated Computing Stack

The Nerve Centre of Ricche

Running GPU-accelerated research infrastructure at scale is like piloting a spacecraft -- hundreds of systems working in concert, petabytes of data flowing through pipelines, and millions of computations executing every second. You cannot operate at this scale without complete, real-time visibility into every layer of the stack.

The Ricche Control Room is our answer. A unified operations dashboard that gives infrastructure operators and research leads a single pane of glass across the entire platform -- from GPU memory allocation to data pipeline health to experiment progress.

94.2%
GPU Utilisation

128
Active Experiments

2,540
Simulation Jobs

99.97%
Uptime (30d)

Core Dashboard Panels

GPU Cluster Command

Real-time GPU utilisation heatmap across all nodes. Per-GPU metrics: memory, temperature, power draw, CUDA kernel throughput. Training job queue with priority scheduling and estimated completion. Thermal throttling detection and automatic workload redistribution. GPU failure prediction and proactive alerting.

Experiment Mission Control

Live tracking of every active ML training experiment. Real-time loss curves, learning rate schedules, and validation metrics. Per-experiment GPU allocation and cost tracking. Side-by-side experiment comparison with automated early stopping recommendations. Historical experiment performance analytics.

Data Pipeline Operations

Feed-by-feed ingestion dashboard with per-exchange latency and throughput. Real-time data quality scoring with drift detection. Automatic gap detection and backfill orchestration. End-to-end pipeline lineage -- trace any feature back to its raw data source. Multi-region feed redundancy status.

Simulation War Room

Live Monte Carlo and agent-based simulation monitoring. Queue depth, GPU allocation, and estimated runtimes per simulation batch. Result visualisation with scenario coverage heatmaps. Pass/fail indicators with statistical confidence levels. Scenario stress-test coverage analysis.

Infrastructure Intelligence

Modern research infrastructure is too complex for reactive monitoring. The Control Room uses predictive analytics to anticipate problems before they impact research workflows -- because a GPU failure during a 12-hour training run is not just an incident, it is lost research time that can never be recovered.

Compute Resource Orchestration

Kubernetes cluster overview with real-time pod scheduling and resource allocation. GPU scheduling queue with multi-tier priority (research, simulation, inference). Predictive auto-scaling based on historical workload patterns. Capacity forecasting with 7-day and 30-day projections. Cost optimisation recommendations.

Storage & Network Intelligence

Distributed storage heatmap: NVMe (hot), SSD (warm), object storage (cold). Automatic data lifecycle management and tier migration. Network throughput between compute and storage with congestion detection. I/O latency tracking (p50, p95, p99). Predictive capacity alerts before you run out.

Model Serving Dashboard (NIM)

NVIDIA NIM microservice metrics: RPS, latency distributions, batch utilisation, and GPU memory per model. Live A/B test traffic routing. Model version deployment history with one-click rollback. SLA compliance tracking with automatic alerting when latency targets are breached.

Incident Command Centre

Active alert dashboard with intelligent severity classification (P1-P4). Automated root cause analysis using infrastructure graph correlation. Mean time to detection (MTTD) and resolution (MTTR) trending. On-call routing with escalation workflows. Post-incident review automation.

Operational Excellence Metrics

What gets measured gets improved. The Control Room tracks a comprehensive set of KPIs designed to maximise research throughput while maintaining infrastructure reliability. These are not vanity metrics -- each one directly impacts our ability to produce better research, faster.

Infrastructure Performance

- GPU cluster utilisation: >90% target (currently 94.2%)
- GPU memory efficiency: >80% target with smart workload packing
- Platform uptime: 99.95% SLA (currently 99.97% over 30 days)
- P99 inference latency: <5ms target (currently <2ms via TensorRT)
- Data pipeline end-to-end latency: <500ms for streaming feeds

Research Productivity

- Concurrent experiments: 128 active GPU experiments
- Experiment iteration speed: significantly faster than CPU-only infrastructure
- Model validation throughput: thousands of simulation scenarios per candidate
- Feature engineering pipeline speed: GPU-accelerated via RAPIDS
- Researcher time-to-first-experiment: <30 minutes for new projects

Operational Reliability

- Mean time to detection (MTTD): <2 minutes for P1 incidents
- Mean time to resolution (MTTR): <15 minutes for infrastructure issues
- Alert signal-to-noise ratio: >85% actionable (minimise fatigue)
- Deployment frequency: multiple per day with zero-downtime GitOps
- Data quality score: >99.5% across all asset classes

Alert Classification

P1 -- Critical (Immediate)

GPU cluster failure, primary data feed outage, storage corruption, or complete service unavailability. War room activated. All hands on deck.

P2 -- High (30 min)

Degraded GPU performance, partial feed disruption, elevated latency, or data quality score drop. On-call engineer engaged with escalation path.

P3 -- Medium (4 hours)

Non-critical feed gap, backup delay, capacity warning, or non-production component degradation. Tracked and scheduled for resolution.

P4 -- Low (Next day)

Informational: certificate renewals, dependency updates, cleanup tasks, and capacity planning triggers. Queued for routine maintenance.

Interested in Our Infrastructure Vision?

The Control Room is being built as the operational backbone of our GPU-accelerated research platform. Get in touch to learn more about our progress and roadmap.

info@ricche.ai | ricche.ai