

General Intuition Case Study

→ general intuition

INDUSTRY
Artificial intelligence

HEADQUARTERS
New York, NY

USE CASES
World model and policy training

 CoreWeave

Explore more AI stories

See how CoreWeave accelerates enterprise AI innovation.

[View case studies](#)

EFFICIENCY AS ESSENTIAL

3

Types of models

+B

of video gameplay data from Medal.tv

Challenge

Used by gamers since 2015, Medal.tv houses one of the world's richest collections of real human gameplay data. Recognizing the value of this data, CEO and co-founder Pim de Witte launched General Intuition, an AI research lab dedicated to developing models capable of spatio-temporal reasoning.

However, training world models, policies, and inverse dynamics models (IDMs) is fundamentally different from training large language models. There are no established scaling laws, standard architectures, or proven infrastructure patterns that help determine how model size, data volume, and distributed training interact.

Working with a lean infrastructure team, General Intuition explored several proof of concepts with different cloud providers for processing massive video datasets and running long-lived, multi-node jobs.

Core Needs List

- **Reliable long-running training:** Support multi-day, distributed training jobs for world models, policies, and IDMs without frequent failures
- **High-throughput video pipelines:** Move and process massive gameplay datasets efficiently
- **Expert-led infrastructure support:** Resolve GPU, networking, and scheduling issues without pulling researchers away from research

Solution

The team set out to determine which provider could reliably scale their model training jobs and remove the infrastructure management burden from their ML researchers. Having worked together for previous GPU capacity, General Intuition ran one of their benchmark evaluations with CoreWeave.

Through CoreWeave ARENA, they benchmarked real training pipelines (not simulations) on production-grade GPU clusters and observed distributed training behavior under load in a production-scale AI lab.

Combined with direct access to expert infrastructure engineers, this structured evaluation environment enabled General Intuition to understand how their jobs performed, cost, and scaled before committing to production.

01 Real workload evaluation at production scale

Through CoreWeave ARENA, General Intuition benchmarked their actual world-model training pipelines—not synthetic tests—on production-grade GPU clusters, allowing the team to observe real distributed training behavior under sustained load.

02 Reliable multi-day distributed training

CoreWeave provided an environment where long-running, multi-node training jobs could run for days without frequent interruptions, compared to multiple failures a day on other providers.

03 Operational simplicity for researchers

Using SLURM on Kubernetes (SUNK), General Intuition ran large distributed training jobs without forcing researchers to manage complex orchestration systems. Jobs started easily, retried automatically when needed, and behaved predictably.

04 Direct access to expert infrastructure support

General Intuition worked directly with CoreWeave infrastructure engineers over Slack to troubleshoot performance issues, GPU behavior, and networking concerns.

05 Visibility into infra health and model behavior

Mission Control's provided baseline visibility into node availability and utilization during large-scale runs, while integrations with Weights & Biases allowed General Intuition to track iterations per second, loss curves, and performance trends across providers.

REAL TEST, REAL INSIGHTS

“

“On other providers, our distributed jobs were failing on a daily basis, sometimes multiple times per day. On CoreWeave, long-running training jobs were far more stable.”

Florian Laurent
Infra Lead / Member of
Founding Team

PERFORMANCE AT SCALE

2
weeks

Time from evaluation to production

20–30%

performance

For multi-node training runs

Outcomes

By evaluating their model workloads through CoreWeave ARENA, General Intuition replaced uncertainty with actual model performance evidence. The team gained confidence in how their training pipelines performed at scale, reduced infrastructure friction for researchers, and moved from proof-of-concept evaluations into production training on CoreWeave without slowing research velocity.

Faster path from evaluation to production

General Intuition was able to move from proof-of-concept testing into production training in weeks, not months. Clear performance and scaling visibility removed guesswork and shortened the time required to commit to large-scale GPU deployments.

Reliable large-scale training for all model types

Multi-node training runs on CoreWeave were 20–30% more performant than workloads on other providers. Fewer interruptions meant less wasted compute and allowed researchers to focus on advancing model capabilities instead of restarting jobs.

From firefighting to forward progress

Real-time access to CoreWeave's infrastructure experts, seamless Slurm on Kubernetes with SUNK, and integration with Weights & Biases meant researchers could focus on research outcomes rather than infrastructure troubleshooting.

Power your AI innovation

Learn how CoreWeave serves as the force multiplier for enterprises and AI natives worldwide.

[Contact Us](#)

**Real workloads.
Real results**

Take a look at
CoreWeave ARENA.

[Read more →](#)

**The Essential Cloud
for AI**

Get unmatched speed,
scale, and performance.

[Read more →](#)

**Introducing
CoreWeave ARENA**

Prove performance in a
production-scale AI lab.

[Read more →](#)