CoreWeave | Miris

Content Distribution

Los Angeles, CA

Use Case
E-commerce, Physical AI, and Media & Entertainment

COREWEAVE X MIRIS

Challenge

Miris, a cutting-edge 3D technology company specializing in real-time 3D streaming and spatial content delivery, saw a significant opportunity quickly evolving. Advances in generative AI were dramatically lowering the barrier to creating high-quality 3D content, unlocking a surge in volumetric data across industries.

But delivering that content at scale remained a challenge. Existing approaches forced a tradeoff between quality and scale: either compressing assets for web delivery or relying on pixel streaming that couldn't efficiently reach large audiences.

To unlock the full potential of this new wave of 3D content, Miris needed an AI cloud platform that could deliver photorealistic experiences in real time, without compromise.

- **Preserve visual fidelity** without compromising quality across devices
- **Scale to global audiences** without per-user GPU dependencies
- **Deliver 3D content faster** with minimal latency and load times
- **Provide cost economics** that scale like video

Solution

Miris partnered with CoreWeave to build a new purpose-built architecture for delivering photorealistic 3D content—one that eliminates the traditional tradeoff between quality and scale.

By combining AI-driven volumetric streaming with CoreWeave's AI cloud, Miris can now process massive 3D datasets upstream and stream them more efficiently to end users.

1 **Upstream GPU processing**
Instead of allocating GPUs per user, Miris processes 3D assets once upstream using CoreWeave infrastructure, dramatically improving scalability and efficiency.

2 **Volumetric streaming architecture**
Miris uses GPU compute upstream to optimize 3D assets into an adaptive spatial streaming format, enabling adaptive, high-fidelity streaming across devices and network conditions.

3 **High-performance AI infrastructure**
CoreWeave's cloud platform, including tightly integrated storage and Kubernetes orchestration, supports massive throughput and parallel processing.

4 **Elastic scaling for asset pipelines**
The platform handles thousands to millions of assets per day, ensuring consistent performance for demanding workloads.

5 **Adaptive streaming delivery**
Content dynamically adjusts to network conditions, enabling instant interaction without buffering or degradation in experience.

1B+ Simultaneous viewers
Global 3D delivery

< 200ms Avg. time to first read
Near-instant streaming

~10× more efficient
Higher GPU VRAM efficiency

> We saw a massive wave of 3D content coming. The challenge was delivering that content at scale without compromising quality or performance.

Will McDonald
Chief Product Officer at Miris

Outcomes

With CoreWeave, Miris delivers photorealistic 3D experiences instantly at global scale. This unlocks new applications across retail, media, robotics, and digital twins while eliminating traditional tradeoffs between quality, cost, and performance.

Global scale without GPU bottlenecks
Instead of allocating GPUs per user, Miris processes 3D assets once upstream using CoreWeave infrastructure, dramatically improving scalability and efficiency.

Near-instant interactivity.
Load times and user drop-off? Solved. Miris delivers interactive 3D content instantly, with adaptive streaming that responds to network conditions. This keeps users engaged and improves conversion and retention.

Higher GPU efficiency, lower cost per workload.
By streaming only what each device needs and optimizing memory usage on both ends of the pipeline, Miris extracts significantly more output from the same GPU footprint. Upstream, asset optimization runs efficiently on CoreWeave infrastructure with high VRAM utilization. Downstream, adaptive delivery means client devices receive only the spatial data required for what the user is viewing, not the full asset. The result is more throughput per GPU and lower total cost of ownership.

> “
>
> The reason why we partner so closely with CoreWeave goes well beyond GPUs. We utilize CoreWeave as a true Cloud AI platform across its compute, storage, and networking that gives Miris a differentiated advantage on scale and performance.
>
> **Will McDonald**
> Chief Product Officer at Miris

Scale confidently on CoreWeave

Run, scale, and optimize next-generation AI workloads with purpose-built infrastructure

Contact us

Explore AI pioneers

See more companies building next-gen AI with CoreWeave

View case studies

Explore AI inference

Solutions that sustain performance under real-world demand

Read more

CoreWeave Acquires Monolith AI
Solving complex engineering problems together, including 3D data

Read more

3D Content in Real Time
Go deeper into how Miris delivers photorealistic experiences

Watch video