

Weights & Biases is the AI developer platform powering the AI industry and is trusted by over 1,500 customers including more than 30 foundation model builders. Teams rely on Weights & Biases to deliver AI agents, applications, and models with confidence.

Weights & Biases provides a comprehensive developer platform to productionize AI. You can pre-train and post-train LLMs for agentic tasks, then evaluate, iterate, monitor, and safeguard agents. Weights & Biases brings together all the tools you need into a single platform with enterprise-grade performance, scale, governance, and security.

“

Weave enables our engineers to score changes to product features using different LLMs under the hood. This allows them to make an informed decision about which one to use, taking into consideration not just quality but also cost.

—Geoff Breemer, Senior MLE, Canva

Faster AI application iteration

















- Rigorous evaluations across accuracy, latency, cost, and more
- Tracing for production monitoring and debugging
- Post-train AI agents using RL for company-specific tasks
- Readily access open-source models hosted by CoreWeave


Safer AI applications

- Guardrails protect users and brand from prompt attacks and harmful output
- Implement a system of record for AI using Registry
- Enterprise grade security and support for regulatory compliance

Flexible AI applications

- Future-proof for scaling AI to more workflows
- Flexible deployment options including cloud and on-prem
- Choice of cloud service providers including AWS, Microsoft, Google Cloud, and CoreWeave

Challenge	Solution
Post-train agents with serverless RL	 Training
Debug AI agents and applications pre-production	 W&B Weave
Run rigorous evaluations of AI agents and applications	 W&B Weave
Prompt engineering playground	 W&B Weave
Observability tools for agentic systems	 W&B Weave
Block prompt attacks and harmful output	 W&B Weave
Collect human feedback and annotations	 W&B Weave
Monitor AI agents and applications in production	 W&B Weave
Hosted open-source models on CoreWeave	 W&B Inference
Experiment tracking for reproducibility and governance	 W&B Registry
System of record for AI	 W&B Registry
Track lineage for datasets, models, and metadata	 W&B Registry
Track experiments for fine-tuning and training AI models	 W&B Models
CoreWeave observability during training	 W&B Models
AI model management	 W&B Models
Communicate findings across stakeholders	 W&B Models



LEARN MORE