



MindWalk Accelerating Biologics Discovery with AI on Vultr Cloud GPUs

MindWalk, formerly ImmunoPrecise Antibodies, uses Vultr Cloud GPUs with AMD MI300X to accelerate biologics discovery with AI-powered protein language models. Powered by the LENSai™ platform and HYFT® technology, MindWalk unifies sequence, structure, function, and literature into a single computational language and closes the loop with an integrated, full-stack wet lab, transforming drug discovery and development.

vultr.com

Accelerating Drug Discovery with AI Models on Vultr Cloud GPUs

AI models are improving healthcare by accelerating research, enhancing workflows, and enabling precise, data-driven treatment, powered by trusted sources like PubMed and clinical data.

The GenAI model landscape is shifting. While large models (LLMs) offer broad capabilities, small language models (SLMs) are gaining traction for domain-specific tasks, delivering high performance at lower cost with faster, more secure deployment. Gartner projects that by 2028, 30% of GenAI workloads will run on SLMs deployed on-premises or on-device, up from under 1% in 2024.¹

How LLMs in healthcare are different:

- **Trusted data sources:** Biomedical LLMs use curated databases to reduce hallucinations and boost accuracy.
- **Extended context handling:** Designed for long-form reasoning across papers and patient records to connect symptoms, diagnostics, and treatments.
- **Cross-modal intelligence:** Combines structured biological data with unstructured text for insights into genes, proteins, diseases, and pathways.
- **Integrated pipelines:** Breaks down silos by linking text, sequences, and structured data, enabling end-to-end biomedical discovery.
- **Stacked LLMs with realignment layers:** Platforms like LENSai use multiple LLMs alongside HYFT[®] technology to align AI outputs with biological context – without training custom models.

Gartner predicts that by 2027, 35% of agentic AI will scale through domain-specific reasoning¹. Open-source [small language models](#) support this shift, offering a sustainable, secure, and cost-effective path to specialized AI.

Common use cases for AI models

Large language models (LLMs) and specialized biological LLMs are now driving real-world healthcare systems. LLMs handle broad knowledge tasks, while protein language models (PLMs) deliver speed and precision for areas like molecular analysis and compliance. MindWalk uses LLMs on Vultr Cloud GPUs to accelerate biologics discovery with high-throughput, cost-effective performance by:

- **Literature mining and knowledge retrieval (LLMs):** Extract drug-target interactions or trial outcomes from scientific literature at scale.
- **Drug discovery (PLMs):** Predict protein function by combining protein language models with biological databases and embeddings.
- **Regulatory intelligence (specialized AI models):** Analyze FDA/EMA documentation for safety profiling, compliance, and evidence generation.

Industry

Healthcare & Life Sciences

Use Case

AI-Powered Biologics Discovery with LLMs and Generative Models

Key Technologies

AMD Instinct™ MI300X, Vultr Cloud GPU, PyTorch, LENSai™, HYFT®, AMD ROCm™ open software

Customer Spotlight

MindWalk, formerly known as ImmunoPrecise Antibodies

“The AMD MI300X on Vultr’s high-performance cloud has allowed us to scale large AI models and pipelines that previously hit memory or cost ceilings. It’s more than a spec sheet, it’s a cornerstone for what comes next in computational biology.”

Frédéric Chabot

Head of Corporate Development, MindWalk

Vultr: Purpose-built cloud for AI in healthcare

- Global scale with industry-leading price-to-performance – meet growing AI demands without overspending.
- HIPAA-ready cloud infrastructure designed for healthcare compliance and data privacy.
- Access to AMD Instinct™ GPUs, including the high-memory 192GB MI300X, ideal for large protein models, biomedical LLMs, and generative applications.
- Start fast – no code changes needed with AMD ROCm and prebuilt PyTorch containers.
- Optimized for AI-native healthcare, from embeddings and regulatory models to foundation models and cross-modal protein intelligence.

MindWalk's AI-native approach to biologics on Vultr Cloud GPU

MindWalk uses LENSai™, an AI-native platform linking sequence, structure, and function for biologics discovery. To support large-scale healthcare AI, MindWalk benchmarked AMD Instinct™ MI300X GPUs on Vultr against NVIDIA HGX H100 across three R&D stages:

Natural Language Processing (NLP) for target identification

Protein embeddings for functional insight

Generative models (RFdiffusion) for de novo protein design

Use case: Protein language models for function prediction

Protein language models are specialized SLMs trained on amino acid sequences to predict protein function accurately. Unlike general LLMs, they capture molecular-level biology, making them ideal for drug discovery and engineering.

Challenge

MindWalk needed to run large protein models like ESM-2 to interpret function, predict mutation effects, and enable faster therapeutic development. Memory bottlenecks and compute costs on legacy hardware limited scalability.

Solution

Using AMD Instinct MI300X GPUs on Vultr, MindWalk deployed ESM-2 models via AMD ROCm™-enabled PyTorch containers without rewriting code. Vultr's infrastructure ensured reproducibility and stability at scale.

ESM-2 Protein language model benchmarks

Model Size	Small (35M)	Medium (650M)	Large (3B)
NVIDIA H100 seq/sec	2,482.41	368.04	111.19
AMD MI300X seq/sec	3,413.15	637.94	178.76
Cost Reduction	~44%	~63%	~55%

Higher throughput: AMD Instinct MI300X achieved up to 63% improvement in inference throughput.

Cost savings: Significant reduction in cost per million protein sequences processed.

Better concurrency: Stable performance under large batch sizes and multithreaded inference.

Integrated platform: LENSai™ powered by HYFT®

LENSai™ combines AI with HYFT® – a patented system of bio-native fingerprints that encode sequence, structure, and function into unified biological patterns. This enables:

- Accurate prediction of mutation impact
- Context-aware comparisons across proteins
- Scalable discovery of therapeutic candidates

Together, they transform raw biological data into actionable insight.

Results

Faster time-to-insight: Accelerated literature mining, target discovery, and protein interpretation.

Reduced costs: AMD Instinct MI300X delivered substantial savings across NLP, protein models, and generative design.

Future-ready stack: Vultr's AMD GPU offering supports MindWalk's scaling roadmap for AI-native discovery.

Additional resources



[Blog | AI-Driven Biology with AMD Instinct™ GPUs on Vultr →](#)

[Case study | Learn more about MindWalk →](#)

[Whitepaper | Agentic AI whitepaper →](#)

Advance AI-driven healthcare with Vultr

Utilize AMD MI300X and MI325X on Vultr Cloud GPU infrastructure for drug discovery and genomics.

Get started at vultr.com

