MINDWALK™ ACCELERATES AI DRUG DISCOVERY WITH AMD

CASE STUDY

MindWalk scales rapid discovery with AMD Instinct™ MI300X GPUs and AMD EPYC™ CPUs on Vultr cloud, running bigger models faster, reducing months to hours

AMD × MindWalk™

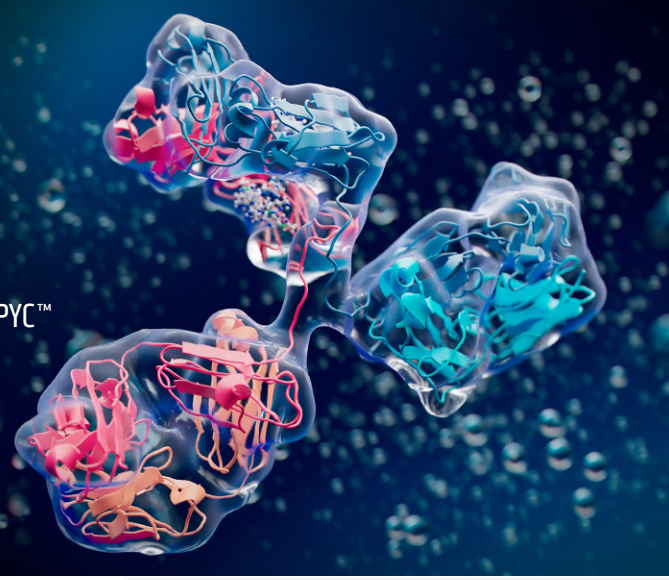

MindWalk™ provides an AI platform that transforms raw biology into connected knowledge, which pharmaceutical scientists use for drug discovery. Its LENSai™ platform is a bio-native knowledge graph powered by HYFT™ technology that links protein sequences, 3D structures, assay results, and peer-reviewed literature into traceable, semantically connected nodes. A HYFTTM is a universal biological fingerprint that captures sequences, 3D structures, lab-measured functions, and peer-reviewed evidence. Instead of chopping data into coarse chunks as conventional large language models (LLMs) do, HYFTTM lets LENSai work at a level biologists care about, including motifs, domains, binding sites, and specific evidence snippets, so teams can identify targets, design and optimize antibodies, and surface off-target or immunogenicity risk earlier. Some examples of what the platform supports are AI-based GLP-1 peptide design and in-silico antibody and vaccine discovery, enabling rapid candidate generation and high-throughput proteome-scale analysis.

> **"The bandwidth on AMD Instinct MI300X GPUs enables larger batches, steadier utilization, and faster convergence across LENSai workflows."**
> Joachim Schreurs PhD, Data Scientist, MindWalk

To run this work at scale, MindWalk uses an integrated architecture built on AMD Instinct™ MI300X GPUs and AMD EPYC™ processors. AMD Instinct MI300X GPUs provide the high memory capacity and bandwidth required to keep large protein language models (PLMs) and broader biological context in GPU memory, which speeds up embedding creation, retrieval, and generative design. AMD EPYC processors deliver CPU throughput for data preparation, scoring, and large parallel screens. Together, they enable MindWalk to continuously update literature and models, shorten the time to insight, and reduce reliance on engineering workarounds that can slow research.

CHANGING THE NATURE OF DISCOVERY

"We realized in 2018 that AI would become a critical part of how we discover and design drugs, but only if the platform could connect all the signals we care about, from literature to lab data, without slowing researchers down," says Frédéric Chabot, Head of Corporate Development, MindWalk.

Dirk Van Hyfte MD PhD, Co-founder, MindWalk, adds, "Classical drug development is very sequential. As a result, teams often uncover critical risks very late, sometimes after years of work. We set out to redesign the workflow, so key checks happen in parallel, compressing timelines and reducing the accumulation of risk."

INDUSTRY
Healthcare and Sciences

CHALLENGES
Keep full biological context in GPU memory so large PLMs run stably and efficiently; embed millions of papers in real time to keep answers current; make always-on drug screening affordable

SOLUTION
An integrated architecture on Vultr cloud featuring AMD Instinct™ MI300X GPUs, AMD EPYC™ processors, and AMD ROCm™ software to fit bigger models, raise throughput, and simplify operations

RESULTS
39% lower cost per million samples in literature mining and ~70% higher embedding throughput on AMD Instinct MI300X GPUs. AMD EPYC CPUs screened 170k antibody pairs in ~4.5 hours vs 145 days

AMD TECHNOLOGY AT A GLANCE
AMD Instinct™ MI300X GPUs
AMD EPYC™ processors
AMD ROCm™ software

TECHNOLOGY PARTNER
VULTR

> **"With 192 GB of high bandwidth memory, AMD Instinct MI300X GPUs let us run even our largest PLMs and batches on one accelerator."**
>
> Joachim Schreurs PhD, Data Scientist, MindWalk

MindWalk sought infrastructure that could solve three problems at once. First, large models and larger batches require high memory capacity and bandwidth to ensure that the full biological context fits and remains stable during execution. Second, retrieval pipelines must embed and index millions of scientific papers and return the right passages in real time so LENSai stays current. Third, the economics must support continuous operation across teams, not occasional one-off jobs. With those needs defined, MindWalk sought a platform that could accommodate its largest models, sustain high-end-to-end throughput, and support high-throughput CPU-based screening at a predictable cost.

BUILDING AN ARCHITECTURE FOR CONTINUOUS DISCOVERY WITH AMD

MindWalk runs parts of its GPU workloads on Vultr, a cloud provider of high-performance infrastructure, and centers the platform on AMD Instinct MI300X GPUs to provide ample room for large biology models to run. "With 192 GB of high bandwidth memory, AMD Instinct MI300X GPUs let us run even our largest PLMs and batches on one accelerator instead of splitting work across GPUs," says Joachim Schreurs, Data Scientist, MindWalk. "AMD Instinct MI300X GPUs eliminate sharding, reduce orchestration overhead, and improve embedding throughput."

Bandwidth matters for day-to-day performance. With 5.3 TB/s of memory bandwidth, AMD Instinct MI300X GPUs keep embeddings, retrieval, and diffusion-based design processes fed, so model runs spend less time waiting on data. "The bandwidth on AMD Instinct MI300X GPUs enables larger batches, steadier utilization, and faster convergence across LENSai workflows," says Joachim Schreurs.

Researchers using LENSai™ move faster from large-model runs to informed design choices with AMD Instinct™ GPUs, keeping biology in memory and relevant studies readily accessible.

Software readiness was just as important. AMD ROCm software, the open GPU computing stack from AMD for accelerated AI and HPC, plays a pivotal role. "PyTorch is the backbone of our model development, so having native support on AMD ROCm meant we could transition core workloads without rebuilding the pipeline," says Dirk Van Hyfte.

MindWalk collaborated with AMD engineers and Brium (now part of AMD) to adapt CUDA®-locked life-science libraries, including Deep Graph Library and SE(3)-Transformers. "Close technical support from AMD helped us remove blockers and move core workloads into production on AMD ROCm quickly," says Dirk Van Hyfte. "The open ROCm stack also makes it easier to integrate new frameworks and models as they emerge, keeping the LENSai graph current."

> **"The speed provided by AMD EPYC processors means customers can surface risks earlier and avoid costly surprises."**
>
> Frédéric Chabot, Head of Corporate Development, MindWalk

Keeping GPUs busy requires a strong CPU foundation. AMD EPYC processors handle data preparation, scoring, and high-throughput CPU-based screening, allowing GPUs to focus on the most compute-intensive AI tasks. These include moving large volumes of sequences through quality checks, conducting immunogenicity evaluations across very large candidate sets, and orchestrating jobs to ensure that AMD Instinct MI300X GPUs receive a steady flow of work. "AMD EPYC processors are the workhorses for large-scale screening and orchestration. GPUs take the deepest learning tasks, while EPYC keeps throughput high across the rest of the pipeline," says Joachim Schreurs.

MEMORY AND SPEED ENHANCE DISCOVERY ECONOMICS WITH AMD

(RAG) for literature mining, AMD Instinct MI300X GPUs delivered approximately 39 percent lower cost per million samples in MindWalk's initial testing, at roughly 3,421 sequences per second compared to 2,741 sequences per second. "Lower cost per query and higher throughput make full-scale sweeps across the biomedical literature routine," says Joachim Schreurs. "That is how we surface rare connections, including off-target interactions and discontinuous epitopes, without delaying researchers."

PLM embeddings show how capacity and speed translate into daily gains. MindWalk reported about a 70 percent increase in embedding throughput on AMD Instinct MI300X GPUs compared to the other platform tested. "Keeping models and context in one AMD Instinct MI300X accelerator lets us push bigger batches through with fewer passes," says Joachim Schreurs, "which shortens end-to-end wall time and helps manage compute cost per embedding, making proteome-scale analysis practical."

Generative protein design benefits in the same way. Training and scoring RFdiffusion campaigns are computationally intensive and sensitive to memory movement. The high memory bandwidth of AMD Instinct MI300X GPUs helps keep these workloads fed, so design iterations proceed quickly. "On AMD Instinct MI300X GPUs, RFdiffusion runs fast and efficiently enough to move quickly from searching to designing at pharma scale," says Dirk Van Hyfte.

> **"We are compressing timelines with AMD. Our pharmaceutical clients can test now, prevent pain downstream, and save significant cost."**
>
> Frédéric Chabot, Head of Corporate Development, MindWalk

In an immunogenicity benchmark used to help customers assess risk early, MindWalk found that a cluster of AMD EPYC processors screened over 170,000 antibody pairs in about 4.5 hours across multiple 64-core nodes. On a previously evaluated platform, MindWalk estimated that the same workload would take 145 days. This shift from months to hours was achieved through a combination of better parallelization, improved data transfer between steps, enhanced CPU utilization, and optimized hardware. "The speed provided by AMD EPYC processors means customers can surface risks earlier and avoid costly surprises," says Frédéric Chabot.

SUSTAINING DISCOVERY VELOCITY FOR THE LONG TERM

MindWalk's roadmap is to make continuous discovery routine, ingest new literature, widen proteome-scale embeddings, and run larger generative campaigns without disrupting production work. "With AMD Instinct MI300X GPUs, AMD EPYC processors, and AMD ROCm, we keep the platform current and turn new ideas into results faster," says Dirk Van Hyfte. "We are compressing timelines with AMD. Our pharmaceutical clients can test now, prevent pain downstream, and save significant cost," adds Frédéric Chabot.

MindWalk relies on the strong CPU foundation from AMD EPYC™ processors, enabling AMD Instinct™ MI300X GPUs to focus on the most compute-intensive AI tasks.

EXPLORE AMD INSTINCT GPUS
Sign up to receive AMD data center communications.

ABOUT MINDWALK™

MindWalk™ is a Bio-Native AI company transforming drug discovery and development. Powered by patented HYFT™ technology and the LENSai™ platform, MindWalk™ unifies sequence, structure, function, and literature into a single computational language and closes the loop with an integrated, full-stack wet lab. The platform supports rapid epitope mapping, de novo molecular design, in silico vaccine exploration, and population-scale biologics analytics that help turn insights into validated candidates at speed. For more information visit mindwalkai.com.

ABOUT AMD

For more than 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies. Billions of people, leading Fortune 500 businesses, and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) website, blog, LinkedIn, and X pages.

DISCLAIMERS

MindWalk™ provides all performance and cost savings claims, but AMD has not independently verified them. A variety of variables impact performance and cost benefits. Results herein are specific to MindWalk™ and may not be typical. GD-181.

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change. It may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between different manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes to the content hereof from time to time without obligation of AMD to notify any person of such revisions or changes. GD-18.

COPYRIGHT NOTICE