VULTR

HEALTHCARE & LIFE SCIENCES AI SUMMIT · SEPTEMBER 17, 2026

OpenFold3 on Vultr

Accelerating Drug Discovery, together with AMD

Duncan Ng

VP, Solutions Engineering, Vultr

VULTR + AMD + MINDWALK

AI in life sciences is moving from pilot to production

Most IT and infrastructure environments aren't ready for what comes next — scalable compute, secure data environments, and infrastructure that doesn't force a rebuild every time a workload grows.

- **Medical Imaging:** Radiology, pathology, and diagnostic AI at scale
- **Genomics:** Sequencing and variant analysis pipelines
- **Clinical AI:** Decision support in regulated data environments
- **Drug Discovery:** Structure prediction, screening, and precision medicine

Protein structure prediction is a perfect stress test

Heavy preprocessing on CPU, heavy inference on GPU — in one workflow.

One complete platform — ***not*** pieces to assemble



Vultr Kubernetes Engine (VKE)

Orchestrates deployment, networking, and HTTPS automatically



Vultr Cloud Compute

CPU resources for preprocessing and MSA workflows



Vultr Cloud GPU

Powered by AMD Instinct™ MI325X GPUs for inference



AMD Inference Microservices (AIM) for OpenFold3

One-click deploy from the VKE Marketplace

Minutes, not manual builds

The Marketplace app automates Kubernetes deployment, networking, HTTPS, and model serving.

Customer spotlight: **MindWalk**

MindWalk™

MindWalk builds AI for **drug discovery and biomolecular interaction research**. It independently validated AMD Inference Microservices for OpenFold3 on Vultr Kubernetes Engine using real research workloads.

- **Reduce infrastructure complexity**

  Let researchers focus on prediction, not building AI infrastructure

- **Validate production workflows**

  End-to-end evaluation with representative healthcare and life sciences workloads

Workloads evaluated

- Protein
- RNA
- DNA
- Ligand
- Antibody-Antigen

Results that ***matter***

3.5–33.8 sec

Structure prediction across complexes from 76 to 1,110 amino acids

~5× faster

Antibody–antigen inference vs. MindWalk's prior environment (60 sec → 11.7 sec)

Selected inference times (no MSA)

Workload	Protein size	Inference time
Protein–DNA	76 AAs	~3.5 sec
Protein–RNA	164 AAs	~5.5 sec
Antibody–antigen	568 AAs	~11.7 sec

Why it matters for life sciences

Less time managing infrastructure. ***More time on the science.***

Open AI Ecosystem

Adopt open-source models within a standardized Kubernetes environment — flexibility and portability, no lock-in

Simplify AI Deployment

Reduce deployment complexity through Vultr Marketplace and VKE, minimizing manual configuration

Integrated CPU + GPU Workflows

CPU handles preprocessing and MSA, GPU handles inference — matched automatically within one environment

Repeatable Enterprise Deployment

The same Kubernetes-native model scales from a single researcher to enterprise production

> “Our focus is accelerating biological discovery, not managing infrastructure. Having a standardized way to deploy OpenFold3 allows us to spend more time evaluating scientific hypotheses and less time assembling and operating complex AI infrastructure”.

Dirk Van Hyfte, MD, PhD

Chief Technology Officer, MindWalk AI

VULTR

Ready to deploy OpenFold3?

A repeatable, Kubernetes-native path from pilot to production — for OpenFold3 and beyond.

Contact us at vultr.com

With thanks to AMD and MindWalk