

NeuralMesh™ Powers Breakthrough Drug Discovery Research

NeuralMesh™ accelerates AI-driven drug discovery with a fully distributed storage architecture that scales to billions of files across hybrid storage environments, cutting pharmaceutical research runs from months to days.

Challenges

- Slow pipelines delay time to insight
- Billions of small files stall jobs
- Data staging steals research hours
- Legacy storage caps research output and starves GPU clusters
- Rising cost per experiment
- Compliance and regulatory data retention requirements

Solution

NeuralMesh™ runs as a distributed, shared-nothing architecture where every node serves data, metadata, and I/O in parallel across a single namespace hosted on premises, in the cloud, or in hybrid settings. NeuralMesh Axon™ pools NVMe that already exists inside compute servers, activating this capacity and saving precious research dollars. Augmented Memory Grid™ extends the key-value cache to speed large-model inference by up to 1,000x.

Benefits

- Accelerate time to discovery
- Cut run times from months to days
- Drive GPU utilization past 90%
- Run more concurrent research jobs
- Scale to billions of files in a single multi-tenant architecture
- Simplify data management
- Federate data across hybrid cloud

Turn faster storage into faster scientific discovery

Pharmaceutical research now runs on AI. Molecular screening, protein structure prediction, and generative chemistry all depend on how fast data moves, and slow storage turns potential discovery into weeks of idle waiting. As datasets grow to billions of small files, storage performance, not the science, sets the pace.

NeuralMesh changes the pace of discovery. Faster storage performance translates directly into breakthrough research and scientific discovery, moving research pipelines from raw instrument data to results without the staging and tuning that slow legacy systems require.

At the Oklahoma Medical Research Foundation, **research job run times dropped by up to 10x**, with long jobs from 70 days to just 7, freeing up scientists to run more concurrent research.

*“Our scientists can now **more easily complete their data analysis without dealing with data infrastructure**. With NeuralMesh, we have so much performance and expandable capacity.”*

— Stuart Glenn, Assistant Director of Infrastructure & Research Computing, OMRF

