

NeuralMesh™ Scales Banking Operations Independently

NeuralMesh™ scales capacity and performance separately, unifies file and object storage, and isolates banking business units on a shared platform, so records remain queryable for years without slowing daily operations.

Challenges

- Retention rules require access to decades of records
- Seasonal swings strain capacity
- Audit lookbacks slow live banking
- Shared infrastructure blurs business-unit boundaries
- Compliance demands segregation

Solution

NeuralMesh™ spreads data and metadata across every node, so no single path dominates how fast records move. **Fast S3 Object Store** runs file and object on a single copy, so transactions, archives, and audit pulls share one namespace. **Data reduction** runs off the write path, shrinking years of records without slowing a single write. **Multitenancy** gives each business unit its own capacity, network context, and policy boundary on a shared cluster. Combined, this means capacity and performance scale separately, so years of records never tax daily operations.

Benefits

- Scales capacity apart from performance
- Absorbs seasonal peak demands
- Serves audit lookbacks without slowing business operations
- Securely isolates business units on shared infrastructure
- Segregates data for compliance

Spend less money on storage and split growth from performance

Banking runs on two clocks. Records must stay queryable for years after the transaction that created them closes, while daily volume spikes hard at month-end, quarter-end, and the holidays. Storage is often overprovisioned by retention and daily banking uses, so when it is sized for peak volume and archive, **it ends up sitting idle for the rest of the year**. That is a massive capital investment that is going unused. An audit lookback that competes with production banking for the same infrastructure is the slowdown banking can least afford.

A drive failure during a quarter-end close can't turn into a multi-day rebuild that locks records an auditor is asking for. **NeuralMesh rebuilds in parallel across the whole cluster** instead of funneling recovery through one path, so a failure stays a contained event, not a delayed report.

One financial services group moved analytics workloads off NAS onto NeuralMesh and reported **10x faster performance** and seamless backup to object storage running alongside production reads. This allowed them to **cut infrastructure costs by 86%** and **reduce database run time by 66%** against local NVMe. This mirrors the same capacity vs. performance tradeoff banking's retention requirements create every day.

"With WEKA, we have so much performance and expandable capacity that we don't have to think about it."

— Stuart Glenn, Assistant Director, Oklahoma Medical Research Foundation