

NeuralMesh™ Unifies Simulation Data, Saturates GPUs for Automotive Manufacturers

NeuralMesh™ absorbs concurrent CFD, FEA, and digital twins workloads, so AV manufacturers keep GPU clusters saturated through vehicle and component simulation and finish design cycles without costly re-spins.

Challenges

- Legacy storage wastes GPU cycles
- Concurrent runs risk restarts when running on legacy storage infrastructure
- Computer-aided engineering (CAE) licenses that cost millions go unused
- Concurrent jobs queue on cluster
- Datasets grow over time and infrastructure costs climb in unison

Solution

NeuralMesh spreads data and metadata across every node, so concurrent CFD and FEA jobs share one GPU cluster, eliminating controller bottlenecks. Parallel I/O keeps each simulation fed at full throughput, no matter how many simulations run at once.

Mid-run checkpoints commit at scale, and real-time rebuilds keep a run performing through failure instead of restarting. A single namespace serves POSIX, NFS, SMB, and S3 together, so CAE tools and archives share one copy. The result is compounding speed, consistency, and accuracy.

Benefits

- Saturate GPUs through simulation
- Protect multi-day CFD/FEA runs
- Free CAE licenses sooner
- Run more jobs on the same cluster
- Cut infrastructure costs by up to 80%
- Drive GPU utilization past 90%
- Stretch usable capacity up to 6x
- Federate data across hybrid cloud

Stop Re-Spins Before They Start in Automotive Manufacturing

An automotive manufacturing re-spin is never a rounding error. It can lead to weeks of Computational Fluid Dynamics (CFD) and Finite Element Analysis (FEA) work that must be redone and a launch date that slips. When a model cannot run the next simulation in a timely manner because data couldn't reach the cluster before the insight went cold, manufacturers lose millions of dollars in productivity losses.

Design, production, and quality teams often share one cluster. They shouldn't have to fight for resources. NeuralMesh's native multitenancy gives each team its own hardware-level isolation and quality-of-service ceiling. That means a peak simulation load never starves the inspection line's real-time GPU feed. With NeuralMesh, no one is blocking progress.

Automotive engineering rarely resides in one place. Design centers and manufacturing plants often span countries or even continents, requiring cloud-burst capacity for peak simulation runs that aren't performed on-premises. NeuralMesh's intelligent replication makes a remote dataset browsable before data has been moved, hydrating only the files a job actually needs. Teams place a CFD run wherever GPU capacity is open instead of waiting for data to catch up. **The payoff: fewer re-spins, tighter launch schedules, and GPUs that stay busy instead of waiting on data.**

"WEKA's storage scalability and ability to grow the infrastructure without losing performance was a key factor in the decision."

INNOVIZ
TECHNOLOGIES

— Oren Ben Ibghei, IT Manager, Innoviz