

Fleet-Level Depth-Selective Training Governance

by [Nick Clark](#) | Published April 25, 2026

What Fleet-Level Training Governance Specifies

The architecture extends Cognition's depth-selective training-governance primitive from cloud-centric federated learning to fleet operation. Operating units (vehicles, drones, robots, infrastructure agents) contribute training data through credentialed observations; the depth-selective gradient routing controls which training contributions affect which model layers; per-example governance-chain provenance traces every gradient update back to the contributing observation.

The fleet-level extension handles patterns specific to operating fleets: contributors that operate intermittently, training data that has spatial-temporal context (where and when the observation was produced), training contributions that may need to be revoked when their authority becomes invalid, and training distribution that propagates through mesh rather than through cloud.

Why Cloud-Centric Federated Learning Doesn't Reach Fleet Operation

Federated learning architectures (Google's Gboard updates, Apple's privacy-preserving learning, the broader academic federated-learning literature) assume cloud-mediated coordination: contributors send gradients to cloud aggregators; cloud

aggregators produce updated models; updated models distribute to contributors. The pattern requires continuous-or-nearly-continuous cloud connectivity.

Operating fleets routinely don't have this. Defense fleets in expeditionary deployment. Maritime fleets between cellular coverage zones. Agricultural fleets in remote regions. Mining fleets underground. Each context faces structural mismatch between cloud-centric federated learning and fleet operating reality.

How Fleet-Level Training Composes With Operation

Each operating unit's training contributions are credentialed observations consumed by the fleet-coordination authority. The authority's depth-selective gradient routing aggregates contributions and produces updated model components; the components distribute through the mesh to other operating units; the cycle continues without continuous cloud connectivity.

Mobile store-and-forward enables fleet-training distribution where mesh connectivity is intermittent. Operating units with current model state can carry it across regions, propagating updates to neighbors when in connectivity range. The fleet's collective model state evolves through structural mesh propagation rather than cloud round-trips.

What This Enables for Edge-Fleet Training

Defense fleets gain depth-selective training governance that operates correctly in expeditionary deployment. Maritime fleets gain federated learning across global routes without satellite-connectivity dependency. Agricultural fleets gain field-level model adaptation without cellular augmentation.

The architecture supports the regulatory pressure that emerging compliance regimes (UNECE R155 cybersecurity, FDA AI/ML SaMD, NHTSA training-data provenance) place on fleet operators. The patent positions the primitive at the layer where fleet operators are increasingly required to govern their training pipelines structurally.