

# Three-Tier Operator Intent Fusion

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

## What the Three Tiers Specify

Tier 1: Full cognitive-state broadcast. Cooperative agents publish their planning graphs, intent fields, and capability envelopes as credentialed observations. Other agents within scope consume the broadcast at full fidelity through governance-credentialed admission.

Tier 2: Structured partial-fidelity bus extraction. Vehicles, drones, or devices publish a limited but specified subset of intent indicators — turn signals, brake lights, route plans, formation orders, transponder data. The partial fidelity is structured enough to be machine-readable and constrained enough to be operationally feasible.

Tier 3: Behavior-inferred attribution. Operating units infer intent from sensor cues — trajectory, gaze direction, gesture, formation, historical pattern. The inference produces a credentialed observation that other consumers can corroborate, contradict, or refine.

## Why Tier-Independent Architectures Produce Brittleness

Current architectures treat the three populations as fundamentally different. Cooperative agents get one protocol (V2X full broadcast); non-cooperative agents get

sensor-based reconstruction (computer vision, motion modeling); adversarial agents get classification heuristics. The three layers don't compose, and the boundary cases produce structural failure.

An autonomous vehicle can coordinate tightly with another autonomous vehicle but treats human drivers as opaque hazards. A drone with cooperative airspace data is paralyzed when a non-cooperative drone enters its window. The architecture has no concept of 'this entity is partially-cooperating' or 'this entity's intent is partially-inferable.' Mixed-fleet operation gets relegated to constrained environments where everyone is fully cooperative.

## **How Cross-Tier Fusion Operates**

All three tiers feed contributions into the same composite admissibility evaluator. The evaluator weights contributions by tier — Tier 1 highest, Tier 2 moderate, Tier 3 lowest — and combines them with environmental observations, governance policy, and operational context to produce a single coherent intent estimate per observed entity.

Tier identification is itself an architectural primitive. An entity may be observed across all three tiers simultaneously: a cooperative vehicle broadcasting Tier 1 also has visible turn signals (Tier 2) and observable trajectory (Tier 3). The evaluator handles agreement and disagreement between tiers structurally, weighting consistent multi-tier observations more heavily and surfacing single-tier divergences for downstream evaluation.

## **What This Enables for Mixed-Fleet Operation**

L4/L5 commercial deployment depends on operating in mixed traffic with human-driven vehicles. The current architecture treats human drivers as Tier 3 only, with the brittleness that produces. Three-tier fusion lets the autonomous vehicle weight Tier 2

indicators (turn signals, brake lights) explicitly while still consuming Tier 3 inference, producing more robust coordination than either alone.

Drone-swarm UTM, defense JADC2, and emerging civil-military airspace coordination all require mixed-population coordination. The same primitive serves all three with configuration changes rather than re-implementation. The patent positions the primitive at the layer mixed-fleet coordination has been missing.