

Three-Tier Intent Fidelity

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

What Three-Tier Intent Specifies

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; adversarial agents get classification heuristics. The three layers don't compose, and the boundary cases produce structural failure.

Three-tier fusion is structural. Each tier feeds the same admissibility evaluator with declared fidelity. Cross-tier corroboration handles agreement and disagreement explicitly. Mixed-fleet operation moves from special-case workaround to architectural primitive.

How Cross-Tier Fusion Operates

The composite admissibility evaluator runs continuously per neighboring entity. For each entity in observation range, the evaluator collects available Tier 1 broadcasts (V2X cooperative messages where present), Tier 2 signals (visual indicators, transponder data, structured partial broadcasts), and Tier 3 inferences (motion-based attribution).

The output is a per-entity intent estimate with declared confidence and fidelity. The autonomous vehicle's planning consumes the estimate as a credentialed observation. Cross-entity coordination operates through the same mechanism for all observed entities regardless of which tier they're contributing through.

What This Enables for Mixed-Fleet Operation

L4/L5 commercial deployment depends on operating in mixed traffic. 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.

