

Waymo's Execution Stack Does Not Externalize Harm Ordering

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

What Waymo's Execution Layer Looks Like Today

Waymo's Driver runs continuous trajectory planning, multi-second forward simulation, and a layered execution stack that gates actuator commands against safety constraints. The stack incorporates ISO 26262 functional safety patterns, an internal Responsibility-Sensitive Safety (RSS) variant, and proprietary heuristics learned across hundreds of millions of operational miles. By the metrics Waymo publishes, the result is one of the safer driving systems on the road.

What the stack does not externalize is the harm ordering it applies when the available actuations all produce some harm. The relative weighting of pedestrian outcomes, occupant outcomes, cyclist outcomes, property damage, and ego-vehicle outcomes is computed inside the Waymo stack, signed by Waymo, and not directly auditable by — or configurable by — the state DOTs and federal regulators whose jurisdiction the vehicle operates under.

Why That's a Structural Liability, Not a Style Choice

The trolley-problem framing has been treated as a philosophical edge case for a decade of AV development. It is in fact a routine engineering decision that every L4

vehicle makes thousands of times per operating mile under sub-second pressure. The vehicle weights occupant safety against pedestrian safety, weights forward-collision risk against rear-collision risk, weights staying-in-lane against avoiding-debris. Every weighting is a harm-ordering decision.

When the harm ordering is hardcoded inside the Waymo stack, the regulatory authority has no structural mechanism to specify it. The authority can review a written ethics statement, can audit logged outcomes after-the-fact, and can revoke operating permission. It cannot configure the ordering. That is a structural mismatch between regulatory authority (which lives at the jurisdiction) and ethical decision authority (which lives at the manufacturer).

How Externalization Would Work Structurally

Confidence-governed actuation specifies harm ordering as a credentialed observation signed by the governing jurisdiction. The state DOT, the NHTSA equivalent, the city traffic authority — each can publish a credentialed harm-ordering policy applicable to its territory. A vehicle entering the territory consumes the credentialed policy through composite admissibility, applies it during execution-mode selection, and records every harm-minimization deviation in lineage with the policy under which it was evaluated.

The architectural change is not from 'no ordering' to 'an ordering.' It is from 'the ordering lives in proprietary Waymo software' to 'the ordering lives in jurisdiction-credentialed governance policy that the Waymo stack consumes as input.' The stack itself can remain Waymo's competitive differentiator; the ordering moves to where the regulatory authority can configure it.

What This Enables for L4 Commercial Deployment

The principal commercial blocker for L4 robotaxi expansion is not technology. It is the liability allocation question: when a Waymo vehicle is in a no-good-options scenario, who is responsible for the choice it makes? Today the answer is Waymo, because Waymo wrote the choice into the software. That answer concentrates liability and produces the regulatory caution that has held L4 to narrow geographies despite a strong safety record.

Externalized harm ordering shifts the allocation. The state DOT signs the ordering; Waymo's stack executes it; the manufacturer is liable for executing the signed ordering correctly, the jurisdiction is liable for the ordering itself. This is the structural pattern that fits how every other regulated transportation domain (aviation, rail, shipping) actually works. Until the AV industry adopts it, L4 stays niche. The patent positions the architecture that enables the adoption.