



[Home](#) [Licensing](#) [Patents](#) [Articles](#)

## Cruise's Safety System Cannot Track Its Own Consistency

by [Nick Clark](#) | Published March 27, 2026 | [PDF](#)

Cruise invested deeply in autonomous vehicle safety, building a framework that includes behavioral safety validation, extensive simulation, and structured incident analysis. But the safety system evaluates each decision against predefined criteria without maintaining persistent state about its own normative consistency. The vehicle does not know whether its cumulative safety decisions form a coherent pattern or whether subtle drift has altered its safety posture. Resolving this requires integrity coherence as a persistent cognitive primitive.

---

### What Cruise built

Cruise developed an autonomous vehicle stack with particular emphasis on urban driving, where the density of interactions between vehicles, pedestrians, cyclists, and infrastructure creates a continuous stream of safety-relevant decisions. The safety framework includes scenario-based validation, statistical

safety metrics, and structured processes for analyzing incidents and near-misses.

The system's safety decisions are governed by cost functions that encode priorities: avoid collisions, maintain safe following distances, yield appropriately, and manage speed in zones with vulnerable road users. Each planning cycle evaluates the current situation against these cost functions and selects the trajectory that best satisfies the safety constraints. The approach is sound and reflects serious engineering commitment to safety.

## The gap between validation and self-awareness

Cruise's safety challenge became publicly visible when incidents revealed a gap between the system's per-decision safety evaluation and its ability to recognize when its overall behavioral pattern had become problematic. An individual decision to continue through an intersection may satisfy all safety constraints. A pattern of decisions that consistently prioritizes traffic flow over pedestrian buffer space may not violate any single rule while representing a normative shift that degrades safety margins over time.

This is the distinction between validation and self-awareness. Validation asks whether this decision satisfies the rules. Self-awareness asks whether the pattern of decisions over the last thousand miles remains consistent with the safety posture the system is supposed to maintain. Current architecture performs validation. It does not perform self-awareness.

The consequences are significant. Post-incident analysis can identify that a series of decisions collectively created an unsafe situation. But the system itself cannot identify that pattern in real time because it has no persistent normative state against which to measure its current behavioral trajectory.

## Why metrics are not normative state

Cruise collects extensive safety metrics: collision rates, near-miss frequencies, safety-critical intervention counts. These metrics measure outcomes. They do not track the internal consistency of the decision-making process that produces those outcomes. A system can maintain acceptable safety metrics while its underlying decision patterns are diverging from their intended normative baseline, until the divergence produces an outcome that the metrics finally capture.

Normative state is different. It tracks the agent's position relative to its own declared values continuously, not through outcomes but through the structure of its decisions. The deviation function computes the distance between what the agent is doing and what its normative trajectory predicts it should be doing. This catches drift before it produces adverse outcomes.

## What integrity coherence enables

With a three-domain integrity model, Cruise's vehicles would maintain persistent state tracking behavioral, normative, and narrative coherence. Behavioral integrity monitors whether actions remain consistent with established patterns. Normative integrity tracks whether the principles governing decisions remain stable. Narrative integrity ensures the vehicle's operational story, the account it would give of why it made each decision, remains coherent over time.

When the deviation function detects that the vehicle's safety margin selections have been gradually decreasing, it triggers a coping intercept before any individual decision violates a rule. The vehicle recognizes that it has been drifting and corrects its trajectory. The correction is structural, not reactive. It happens because the integrity primitive detected normative divergence, not because an incident forced a review.

For fleet operations, integrity state enables cross-vehicle normative comparison. If one vehicle's safety posture has diverged from the fleet baseline, the deviation is detectable and correctable before it produces an incident. The fleet maintains collective normative consistency through individual vehicles' self-monitoring.

## The structural requirement

Cruise's safety framework evaluates decisions rigorously. The structural gap is in self-monitoring: the ability to track whether the cumulative pattern of safe decisions remains normatively consistent over time. This requires integrity as a persistent cognitive state with deviation tracking, the coherence trifecta for self-correction, and coping intercepts that engage before normative drift produces adverse outcomes. The vehicle that monitors its own ethical consistency is structurally safer than one that only validates individual decisions.

[Integrity & Coherence All 21 steps →](#)

Track normative consistency. Detect deviation. Self-correct.

Primary Technical Disclosure

[◦ The Coherence Trifecta: Empathy, Integrity, and Self-Esteem as a Unified Control Loop](#)

Secondary Technical

[◦ Coping Under Empathic Pressure: HSP, Narcissism, and Psychopathy as Control-Loop Intercepts](#)[◦ Three-Domain Integrity Model](#)[◦ Deviation Function  \$D=\(N-T\)/\(ExS\)\$](#) [◦ Self-Esteem as Internal Validator](#)[◦ Deviation as Deterministic Semantic Mutation](#)[◦ Integrity Structural Placement](#)[◦ Empathy as Distributed Moral Load](#)[◦ Coherence Trifecta Control Loop](#)[◦ Coping Intercept Patterns](#)[◦ Integrity Deviation Logging](#)[◦ Integrity Collapse Detection](#)[◦ Redemption Engine](#)[◦ Moral Trajectory Forecasting](#)[◦ Integrity-Aware Trust Slope Validation](#)[◦ Integrity-Confidence Cross-Primitive Coupling](#)[◦ Integrity-Modulated Discovery Traversal](#)[◦ Integrity-Aware Multi-Agent Negotiation](#)[◦ Biological Signal Coupling for Integrity](#)[◦ Policy-Based Integrity Constraints](#)[◦ Integrity Field Portability](#)[◦ Predictive Deviation Alerting](#)[◦ Governed Forgetting](#)[◦ Predictive Social Modeling](#)

Applications (General)

[◦ Autonomous Vehicle Ethical Decision-Making Through Computable Integrity](#)[◦ Financial Trading Systems That Track Their Own Normative Consistency](#)[◦ Integrity and Coherence for Legal Advisory Agents](#)[◦ Integrity and Coherence for Government Policy Agents](#)[◦ Integrity and Coherence for Journalism Editorial Agents](#)[◦ Integrity and Coherence for Environmental Compliance Agents](#)[◦ Integrity and Coherence for Insurance Underwriting Agents](#)[◦ Integrity and Coherence for Social Media Moderation Agents](#)

Applications (Specific)

[◦ Waymo's Ethical Decisions Have No Normative Memory](#) • [Cruise's Safety System Cannot Track Its Own Consistency](#) ◦ [JPMorgan's Trading Compliance Has No Normative Trajectory](#) ◦ [Palantir's Analytics Cannot Monitor Their Own Normative Drift](#) ◦ [Aurora's Self-Driving Stack Has No Normative Memory](#) ◦ [Nuro's Delivery Robots Optimize Without Normative Tracking](#) ◦ [Zoox Plans Maneuvers Without Tracking Normative Drift](#) ◦ [Motional Validates Safety Without Governing Normative Trajectory](#) ◦ [Argo AI's Shutdown Reveals the Cost of Missing Normative Architecture](#) ◦ [comma.ai Learns to Drive Without Learning Ethics](#)  
[Integrity & Coherence overview →](#)

AQ  
deterministic  
autonomy

Legal

Subject to one or more pending U.S. and international patent applications, see [Patents](#) for the current list and status. No license, express or implied, is granted. Any use requires a separate written agreement—see [Licensing](#). Patent applications referenced on this site are pending. Claim scope, if any, is subject to examination and may issue in altered form or not at all. See [Legal](#) for terms and conditions.

Adaptive Query™ is a trademark of Nicholas Clark. U.S. federal registration is pending, federal registration. AQ™, AQ Inside™, Adaptive Index™, Adaptive Network™, Semantic Agent™, @AQ™, AQID™, and Adaptive Coin™ are used as trademarks in connection with the Adaptive Query platform and brand. Other names may be trademarks of their respective owners.

Platform operated by Adaptive Query LLC, which provides patent and trademark licensing services. Copyright © 2025-2026 Nicholas Clark. All rights reserved.

Last updated: 2026-03-03



- [Inventive Steps](#)
- [Licensing](#)
- [Patents](#)
- [Articles](#)
- [Legal](#)
- [Opportunities](#)
- [Sitemap](#)



- 
- [nick@qu3ry.net](mailto:nick@qu3ry.net)
- 72 28 14 36 01



[Invented by Nick Clark](#) | Founding Investors: Devin Wilkie