



AxiLayer AI, Inc. | Independent AI Certification, Auditing and Compliance

Autonomous Vehicles 2027: Four Regulators, Four Different Rulebooks

AxiLayer replaces the annual audit snapshot, not the auditor: the human auditor still signs every finding, but the evidence is Always On or On Demand, current the day it is read.

1. The Convergence

Autonomous Vehicles in 2027: Everyone Agrees on the Benchmark. Nobody Agrees on the Proof. Four regulators, one vehicle, and four completely different ideas of what counts as evidence. Something quietly remarkable happened in automated driving regulation this year, and it has attracted almost no attention. Three of the world's major regulators have converged on the same safety benchmark, in almost the same words. The United Kingdom's statement of safety principles sets the standard at a level equivalent to that of careful and competent human drivers. China's new mandatory national standard requires an automated driving system to be at least as safe as a competent and attentive driver performing the dynamic driving task. The European Union's type-approval regulation tells manufacturers to build their safety argument by reference to the performance of competently and carefully driven manual vehicles. Careful and competent. Competent and attentive. Competently and carefully driven. Three regimes, written independently, on three continents, arriving at the same sentence. And then they diverge completely on the only question that actually decides whether your vehicle can be sold, operated or insured, which is what evidence proves you meet it.

2. Europe's System

Europe wants a system, audited. Commission Implementing Regulation European Union 2022 slash 1426 governs type-approval of the automated driving system of fully automated vehicles, and it was amended with effect from 24 March 2026. Its center of gravity is not a test. It is Article 5, which requires the manufacturer to demonstrate to the type-approval authority, in terms of a safety management system, that effective processes, methodologies, training and tools are in place. That system has teeth. The manufacturer must run periodic independent internal process audits, on the order of every two years. It must push safety management obligations down its supply chain by contract. It receives a certificate of compliance for its



safety management system that is valid for no more than three years and can be withdrawn where major non-conformities are found. Article 9 then specifies what the vehicle itself must record when something happens: activation and deactivation, initiation of emergency operation, system failures, minimal risk maneuvers and collisions, each with the reason for the occurrence, the date, the GPS position, a timestamp accurate to one second, and the software identification. Read that last item again. Software identification, on every recorded occurrence. Europe has decided that the question eighteen months from now will not be whether your system was safe. It will be which version of it was in that vehicle on that day.

3. The UK Entity

The United Kingdom wants an entity, accountable. The Automated Vehicles Act 2024 does something no other regime does. It creates a named legal person, the Authorised Self-Driving Entity, using the Act's own spelling, that carries the duty for how the vehicle behaves. The statement of safety principles consultation closed on 9 September 2026 and sets out ten principles covering traffic law compliance, hazard prediction, behavior in unexpected events, interaction with vulnerable road users and with emergency services, and operation within authorized design limits. The enforcement model is what matters. The Secretary of State must monitor the performance of authorized self-driving vehicles annually, assess whether that performance remains consistent with the safety principles, publish yearly reports on how the fleet is doing, and act where it is not. Authorization is not a gate you pass. It is a status you hold, reviewed every year against evidence somebody has to produce.

4. China's Dossier

China wants a dossier, tested by a third party. GB 44721-2026, the safety requirements for automated driving systems of intelligent connected vehicles, was approved by the State Administration for Market Regulation and the National Standardization Administration on 30 July 2026 and becomes effective on 1 July 2027. It is mandatory, it covers L3 and L4 systems on passenger and commercial vehicles, and it obliges manufacturers to safeguard across four dimensions: safety policy, risk management, assurance and improvement. One element of the Chinese standard has no direct counterpart in the others. Third-party testing bodies conduct confirmatory tests, combining test-track, road and simulation



methods with a review of corporate capability and the safety dossier. Not only the vehicle. The organization behind it.

5. America Unblocked

The United States wants deployment, unblocked. On 3 September 2026 the Department of Transportation released AV 5.0, America Leads, its national strategy for automated vehicles for fiscal years 2026 to 2030. It supersedes every federal AV policy statement going back to 2016. It acknowledges openly that the Federal Motor Vehicle Safety Standards were written on the assumption of a human driver, commits to modernizing them for purpose-built vehicles without steering wheels or pedals, and widens exemption eligibility to domestically produced vehicles rather than mainly importers. It is a strategy, not a rulemaking. It creates no binding obligations. And through the first half of 2026 NHTSA moved steadily to take human-driver assumptions out of the standards, proposing on 16 March to exempt vehicles without manual controls from the transmission shift, defrosting and wiping standards, and on 26 June to remove the hand and foot brake control requirements for purpose-built autonomous vehicles while keeping the stopping distance performance. If you read only those headlines, you would conclude the American burden is falling.

6. The Evidence Moves

Fewer requirements. The same evidence. On 26 June 2026, NHTSA withdrew AV STEP. The ADS-equipped vehicle safety, transparency and evaluation program would have been the American answer to the European and Chinese approach. Participating manufacturers, developers and fleet operators would have submitted application information, undergone independent safety assessments, filed quarterly operational reports and accepted public disclosure of some of the resulting data. The reason it was withdrawn is worth reading closely. Industry comments indicated that participation would have been minimal because the program was seen as too burdensome. Safety advocates argued in the opposite direction, that a voluntary program was not a sufficiently comprehensive regulatory tool. With support from neither side, the program did not proceed. So the structured, voluntary, forward-looking route by which an American operator could have demonstrated safety on its own terms, with independent assessment and a reporting rhythm it helped design, no longer exists. None of this means this approach is wrong. The case for it is real and worth stating plainly. The



Federal Motor Vehicle Safety Standards were genuinely written on the assumption of a human driver. Purpose-built vehicles were held up by requirements about windshield wiping and hand-operated brake controls that have no meaning without one. Faster and more predictable exemptions do unblock deployment, and a regulator that moves quickly on that is doing something useful. The argument here is not that one instrument is better than another. It is narrower and it applies everywhere. Whichever instrument a regulator chooses, the operator's evidence obligation does not disappear. It only changes when it falls due, and to whom.

7. The Live Mechanism

What remains in force is what NHTSA states directly. Defect investigation, recall authority and crash reporting obligations, all fully in force. Look at the difference in posture. A program you enroll in is prospective, cooperative and shaped in advance. An investigation is retrospective, and its scope is set by whoever opens it. On 6 May 2026 NHTSA opened a preliminary evaluation into an automated driving system following crash reports in Texas. That is the live mechanism now. Removing the program did not lower the evidence burden. It moved the moment at which the evidence is demanded, from before deployment to after an incident, and it changed who sets the terms. Meanwhile the agency has said it is developing objective, minimum performance standards for ADS competency through an advance notice of proposed rulemaking, and Congress has three vehicles moving at once: the BUILD America 250 Act approved by the House Transportation Committee in late May 2026, which would set federal safety requirements for commercial automated vehicles with state preemption, the SELF DRIVE Act of 2026, and the Motor Vehicle Modernization Act of 2026. The American benchmark is coming. It is simply not here yet, and the interval does not reduce what you will be asked to show when an investigator calls.

8. The Fifth Regime

And a fifth regime arrives underneath all of them. Where a product requires third-party conformity assessment under listed Union harmonization legislation, and EU vehicle type-approval legislation is on that list, the EU AI Act treats the AI inside it as high-risk through Annex I. That brings the full obligation set: a risk management system, data and data governance, technical documentation, automatic logging, human oversight, accuracy and robustness, a quality management system and post-market monitoring. The



Annex I route applies from 2 August 2028. Teams planning only for type-approval are planning for part of it.

9. The Shared Core

What all four actually require, underneath the differences. Here is the thing that should reframe how a mobility company budgets for this. The four regimes disagree about who assesses, when, and against what document. They agree, without coordinating, on what has to exist. A defined operational design domain, argued rather than asserted. Every one of these regimes turns on whether the system was operating inside the limits it was authorized for. Europe requires operation within the declared domain. The UK's ninth principle is operation only within authorized design limits. China requires defined state indicators. If your operational design domain is a paragraph in a slide deck rather than a testable specification with evidence at its edges, you do not have a domain. You have an intention.

10. Boundary Behavior

Evidence of behavior at the boundary, not in the center. Object and event detection and response under degraded perception. Takeover request timing and driver readiness monitoring for L3. Minimal risk maneuver initiation and completion. Behavior around emergency services, cyclists and pedestrians. Performance in the weather and lighting your fleet actually encounters, not the weather your validation set contains. Europe expects hazard analysis and risk assessment supported by failure mode and effect analysis, fault tree analysis or system-theoretic process analysis. It also says plainly that simulation is not a substitute for physical tests.

11. The Record

An occurrence record that can be read by a stranger. Activation, deactivation, emergency operation, failures, minimal risk maneuvers, collisions. With reason, position, timestamp and software version. This is not telemetry for your engineers. It is a record intended to be read, later, by someone who does not trust you and does not have to. A change history that survives an outside reading. Every over-the-air update, every model version, every parameter change, with dates, the fleet segments affected and a named approver. Software identification is written into the European occurrence record for a reason. The regulator has already anticipated the



argument where a manufacturer says the version that caused the incident is not the version under review.

12. The Organization

An organization that can be audited, not just a vehicle that can be tested. Europe audits the safety management system. China reviews corporate capability alongside the safety dossier. The UK holds a named entity accountable annually. In three of the four regimes, the company is in scope, not only the product. A named owner of the clock. When an incident occurs, someone has to decide within hours what gets reported, to whom, in which jurisdictions, and on what basis. A fleet operating in Europe, Great Britain, China and the United States is subject to four different reporting mechanisms with four different triggers. That is a person, not a policy document. That is an evidence discipline. It is not what most automated driving organizations are instrumented to produce, and it cannot be assembled retrospectively in the three weeks after an investigator calls.

13. Why AxiLayer

This is the gap we built the company for. AxiLayer AI is an independent third-party AI assurance firm. The independence is a structural property of how the company is built, not a sentence in an engagement letter. It is also the reason we are useful in this particular sector. The organizations best placed to help an automated driving program through the next eighteen months are frequently the same organizations that supplied the perception stack, wrote the safety case, or sold the simulation environment the case was validated in. That is tolerable when the reader of the file is a private customer. It is not when the reader is a type-approval authority, a market surveillance authority, a licensing regulator, a plaintiff's expert or an underwriter deciding whether to bind cover on a fleet.

14. The Certificate

What we do is make the evidence exist and keep it current, so that a formal assessment is shorter and first-time, and so that in between, when a regulator or an OEM customer asks what changed, on which units and on what date, there is an answer rather than a search. An assessment engagement ends in a signed AxiLayer AI certificate of conformity assessment, issued on our own authority and backed by a re-verifiable



evidence chain. It states what was assessed, against which frameworks, with what result, and on what date.

15. AxiSentinel

AxiSentinel: how the record stays current. Making the evidence exist once is an engagement. Keeping it current through eighteen months of fleet updates is a platform problem, and it is why we built our own. AxiSentinel encodes regulatory instruments as machine-executable packages and evaluates a client's AI systems against every framework that applies in the markets where those systems operate. It runs continuously for clients who want always-on coverage, or on the schedule the client sets. It is live and deployed today across cloud, on premises and air-gapped environments, which matters here because a great deal of automated driving evidence originates on networks that will never touch the public internet.

16. The Human Gate

131 and more regulatory frameworks are encoded and live, across 45 and more jurisdictions and 22 regulated industries, with autonomous vehicles and mobility among them. For this sector that means the EU type-approval route, the AI Act, the UK regime, the Chinese standard and the United States reporting obligations evaluated together, rather than as five separate projects run by five separate teams who never compare notes. The part that matters most is what the platform does not do. A threshold breach raises a Provisional Alert, which is a flag for human review and not a finding. A qualified AxiLayer AI auditor reviews the underlying evidence, validates or overrides the alert, and applies a cryptographic signature. Only then does it enter the tamper-evident evidence chain and update the Compliance Passport, a live, re-verifiable record of what was evaluated, against what, with what result, on what date. And the platform observes and evaluates. It does not block, it does not remediate, and it does not write into a client's system. An independent assurance body cannot touch what it inspects, and building that constraint into the software is the only way to keep the independence real once the software is running inside a customer's environment.

17. The Next Four Weeks

What to do next. If you build, operate or insure automated driving systems, three things are worth doing in the next four weeks, and none of them



require hiring anyone. One. Write your operational design domain as a testable specification. Not a description. A specification, with the boundary conditions named and the evidence at each boundary identified, including the evidence you do not yet have. Every one of the four regimes turns on this document, and it is the one most programs have never written down in a form an outsider could check. Two. Name the owner of the incident clock. One person who knows, for every market you operate in, what triggers a reporting obligation, to whom, within what period, and what the first hour looks like. If you operate in more than one of Europe, Great Britain, China and the United States, that person needs a matrix, not a memory. Three. Ask what you would hand over on Monday. Pick your highest-exposure deployment and ask your team what they could put in front of an assessor or an investigator next week. Not what they could build. What exists today, retrievable, with software versions attached.

18. The Readiness Gap

The gap between those two answers is your actual readiness, and it is almost always wider than most teams expect. If it would help to work through that third question with someone independent, we run a scoped readiness assessment that produces a written operational design domain classification, a gap list your own engineers can act on, and a view of what the evidence record has to contain before an assessment or an investigation starts.

Comment below, or reach us at contactus@axilayerai.com. No commitment beyond the conversation, and we will tell you honestly if you do not need us yet. Our autonomous vehicles and mobility work is here:

<https://www.axilayerai.com/industries/autonomous-vehicles/>

19. The Proof

Three regulators have now told you, in nearly identical words, that your system must drive as well as a careful and competent human being. None of them has told you what a careful and competent human being's evidence file looks like, because human drivers do not have one. That file is yours to build, and it takes longer than the time remaining. This content is for educational and informational purposes only and does not constitute legal, regulatory, or compliance advice.