ZOOX

Zoox Safety Case Framework

Our safety case comprehensively assesses how hardware, software, and operations integrate and interact to deliver a safer autonomous ride.

Our approach to safety at Zoox

At Zoox, safety is foundational to our mission and the guiding principle for every aspect of our robotaxi development. We design and manufacture our robotaxis from the ground up, developing our own driving software and operating our own fleet. This vertical integration enables direct and continuous optimization for safety from the earliest stages of system architecture and engineering design through every ride taken on public roads.

The Zoox safety case quantifies the safety of our robotaxis within the target Operational Design Domain (ODD), the specific set of conditions under which the Zoox robotaxi is designed to operate safely. Our target is to be significantly safer than a human driver. To achieve this goal, the safety case utilizes a system safety approach to address random and systemic causes of faults and design errors across hardware, software, and operations. Total residual safety risk is then aggregated using a quantitative risk assessment framework.

> Zoox robotaxi in San Francisco.

Systems safety process

Our approach to safety is rooted in systems engineering and follows the system safety process. While traditional component-level analysis focuses on individual part failures, our safety case at Zoox comprehensively assesses how hardware, software, and operations integrate and interact. By approaching safety as a property of the entire system, we ensure safety is built into the overall system design and architecture from the start.

The system safety process begins with a combination of structured hazard analysis and historical learnings to identify, classify, and mitigate risks. We use hazard analysis methodologies such as System-Theoretic Process Analysis (STPA), Failure Modes and Effects Analysis (FMEA), Fault Tree Analysis (FTA), and Hazard and Risk Analysis (HARA), commonly recommended in safety standards and best practices. These hazard analyses drive critical architectural decisions such as redundancy and fault-monitoring strategies, while also directing the development of safety requirements and test scenarios.

These requirements are then implemented and rigorously validated through targeted verification and validation tests spanning hardware, software, and operations. Importantly, the system safety process incorporates learnings from real-world operations, feeding them back into the systems engineering process to drive continuous safety improvement.

Quantitative risk assessment framework

We use a quantitative risk assessment framework to evaluate our safety case across three domains within a target ODD: autonomy behavior safety, robot platform safety, and operational safety.

Safety clearance is the formal gate that brings these three domains together. Before any safety-relevant software release, hardware change, or operations revision, we update the safety case and confirm that the combined, quantified risk across all three domains meets our safety targets. This combined risk is measured via the primary metric of our safety case: the estimated rate of potential Collision, Injury, and Fatality events (CIF), measured in miles per event. This estimate draws on both simulation and real-world driving data. More importantly, the CIF metric must meet or exceed our targets for roads. The sections that follow detail how we quantify and validate each domain's contribution to this estimate.

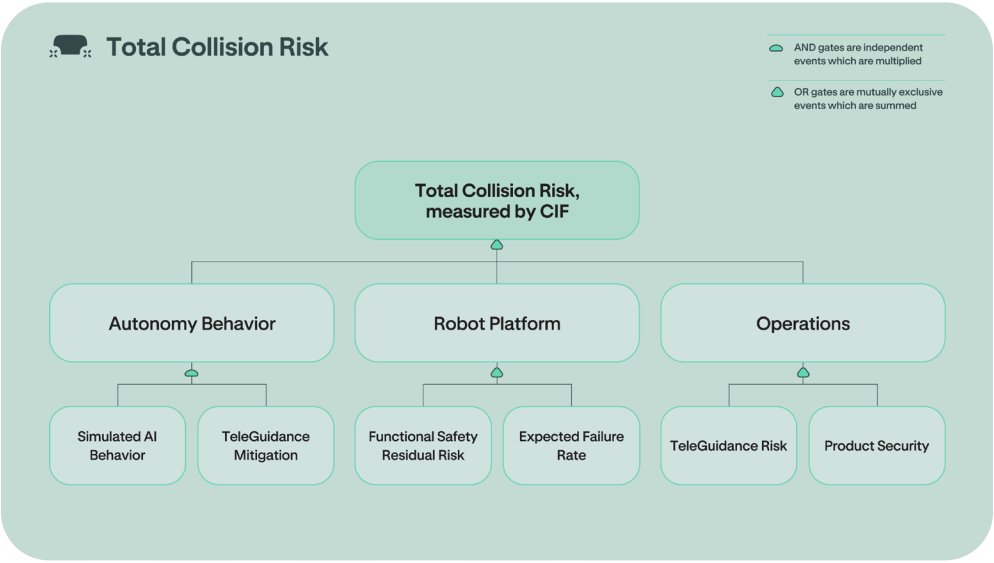

> Image indicating how Total Collision Risk is calculated.

Comprehensive methodologies for quantifying risk

This section first describes Zoox's safety risk metrics. Then, we delineate how risks are assessed in order to demonstrate passing the safety risk metrics.

Our Safety Risk Metrics

CIF is our primary safety risk metric. Similar to other risk frameworks, Zoox assesses safety risk along two dimensions: severity and likelihood. The severity dimension breaks out into three levels: Collision, Injury (MAIS1+[1]), and Fatality. The likelihood dimension (CIF) captures how often events are predicted to occur at each severity level, measured by miles per event (see image below). The CIF metric accounts for harm to any party involved in a collision, including Zoox vehicle occupants, vulnerable road users (VRUs), such as pedestrians, cyclists, and motorcyclists, and the occupants of other vehicles.

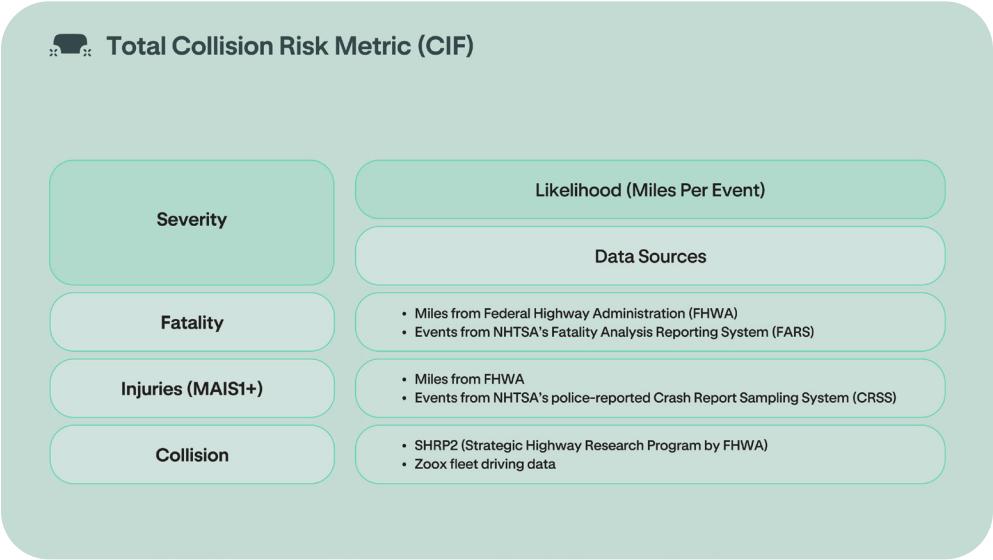

> Image indicating Total Collision Risk is then compared to collision risk human benchmark based on highly credible public collision data.

1. MAIS1+ is defined by the Association for the Advancement of Automotive Medicine (AAAM) as the Maximum Abbreviated Injury Scale, level 1 or higher.

The CIF safety targets are determined by comparison to published human driving collision data (i.e., human benchmark). Zoox sets our safety target to be significantly safer than the human benchmark. This means that our robotaxis need to be significantly safer than human drivers before they are cleared to operate on public roads. This target-setting approach ensures the company's internal safety goals are grounded in real-world human data, customized to the target ODD, and reflect our ambition to be meaningfully safer than human drivers.

Building the human benchmark works as follows. We evaluate credible and published historic human driving collision data, including mileage and event data from the National Highway Traffic Safety Administration's (NHTSA) Crash Report Sampling System (CRSS), NHTSA's Fatality Analysis Reporting System (FARS), the Federal Highway Administration's (FHWA) Strategic Highway Research Program (SHRP2), and FHWA's annual vehicle miles traveled estimates. We then parse this data to account for the difference in risk among varying road speeds (e.g., 25 mph vs. 55 mph) and the mix of road-speed segments within our target ODD. Logged mileage and events from the Zoox test fleet are also used to validate the human benchmark.

The safety performance of the overall system is evaluated based on: the defined ODD, the ADS performance, the safety of the robot platform, and the safety of our operations. To pass our safety clearance, the comprehensive system CIF estimate must meet or exceed our established targets.

We recognize, however, that a pure severity/likelihood-based risk metric is not sufficient to cover rare collision avoidance scenarios, and it is not a replacement for following established industry safety standards, like ISO 26262. Accordingly, we have defined additional safety metrics in other categories. For example, we have a test set to represent rare collision avoidance scenarios [2]. To be considered passing, the Zoox robotaxi must perform on par with or better than nominal human drivers in the simulation. We also follow the ISO 26262 process and other industry best practices to define internal safety requirements for the robot platform systems (e.g., sensors, compute, actuator controllers, and associated firmware). We also define and track other key safety metrics, including rules-of-the-road, near-miss, and operational safety metrics. Adherence to these requirements is also part of the overall safety clearance effort.

The next section explains how we quantify and validate each domain's (referenced in the Total Collision Risk image) contribution to the overall safety case.

2. For more information on edge case engineering, additional information available here: https://zoox.com/journal/edge-case-testing-zoox

Autonomy behavior safety

ADS performance assesses the capability of our autonomy in perceiving the environment, predicting the behavior of other road users, and planning safe trajectories that follow the rules of the road. The ADS accomplishes this through five integrated software functions: localization, perception, prediction, planning, and control. We also have an extra layer of collision checking, which acts as an independent check of the primary system's planned trajectory and, using its own perception and algorithms, determines whether that trajectory is safe to execute or whether it should intervene.

Our safety case evaluates the performance of the autonomy stack, specifically within our defined ODD. By rigorously testing our software against the human benchmark, we validate that the robotaxi can navigate these complex environments with a level of precision and caution that minimizes safety risk. To validate the safety case at the required level of rigor, Zoox employs a comprehensive set of complementary, yet distinct, methodologies.

Zoox employs simulation-based measurement methodologies that include both synthetic simulation scenarios and real-world log-based simulation. Structured hazard analysis and historical safety data inform the coverage of the synthetic simulation pipelines. These pipelines use optimization techniques to search across driving scenarios to discover the conditions where a simulated collision is most likely to occur. Testing further narrows in on those high-risk scenarios by generating refined variationsof the test scenario. Simulation results are then weighted by real-world fleet exposure data to produce statistically grounded safety risk estimates. Complementing synthetic simulations, log-based simulation replays actual fleet driving data through the latest software stack using machine-learning-based sampling to prioritize rare and safety-critical events. Together, these approaches provide overlapping, yet distinct, perspectives on driving software safety.

Additionally, we apply targeted structured testing methodologies for scenarios driving logs that are less likely to surface. One example is using closed-course structured testing to stage perception-sensitive scenarios at controlled test facilities. These scenarios are difficult to reproduce accurately in synthetic simulation and occur infrequently in driving logs, so testing them directly on a closed course lets us evaluate them safely and repeatedly. As mentioned earlier, our dedicated collision avoidance testing methodology targets simulation scenarios that may have a low quantitative impact on CIF because they occur extremely rarely and are difficult to find in driving logs, but represent high-collision avoidance situations, especially involving VRUs.

After all autonomy behavior methodologies have been applied, we validate the safety case on the road, driving and qualifying the ready-to-be-cleared software in our retrofitted test vehicles, with human drivers monitoring vehicle behavior, before driverless clearance is approved.

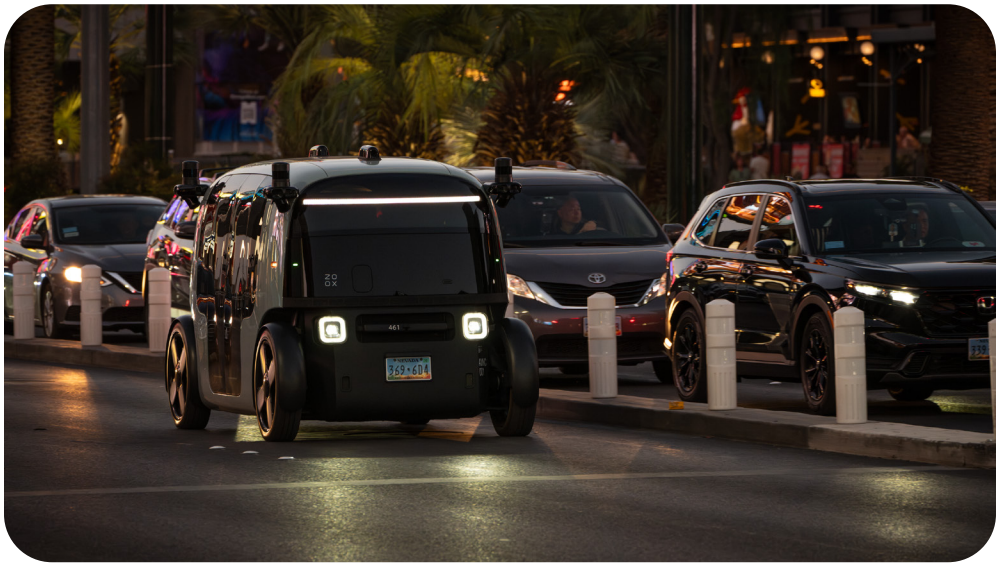

> Zoox robotaxi driving in Las Vegas.

Robot platform safety

Zoox designs its own robotaxis from the ground up, and safety is built directly into the architecture rather than retrofitted into a traditional vehicle design.

To develop safety requirements for the robot platform, Zoox follows the ISO 26262 functional safety process. This begins with a structured hazard analysis, HARA, that identifies what could go wrong with the platform and how serious each hazard would be. From the HARA, we define safety goals and assign each an Automotive Safety Integrity Level (ASIL), a rating of how critical it is to safety. We then capture those goals in a functional safety concept and technical safety concept, which translate them into concrete safety requirements across sensors, compute, actuator controllers, and firmware. We then apply a set of complementary analyses to evaluate the design against those requirements. Failure Modes, Effects, and Diagnostic Analysis (FMEDA) quantifies the hardware architectural metrics and evaluates whether ASIL requirements are achieved, given known hardware faults and the diagnostic coverage. FTA evaluates how well redundancy and onboard safety monitors cover single-point faults. Finally, dedicated fail-operational and fail-safe analyses define the steps the platform takes to reach a safe state once a fault occurs.

The robot platform's requirements are then verified and validated during safety clearance through a mix of software-in-the-loop and hardware-in-the-loop fault-injection testing and closed-course on-vehicle testing. Grounded in the ISO 26262 process, these methodologies produce quantitative estimates that feed directly into the overall safety risk model, ensuring that the physical platform can gracefully mitigate risk and protect occupants in the event of a fault. The robotaxi design includes hardware redundancies to ensure a level of functionality of safety-critical systems even after a fault has occurred. The fail-operational and fail-safe strategies are designed to optimize for both safety and mission execution.

Additionally, the robotaxi is tested and verified to meet or exceed the performance requirements in the applicable Federal Motor Vehicle Safety Standards (FMVSS).[3]

3. While Zoox is exempted from certain FMVSS, we nevertheless design and test our systems to assure that they meet the safety purpose of those requirements (e.g., we approach our sensor clearing system like a conventional windshield wiping system).

Operational safety

The Zoox operational safety program includes continuous and critical monitoring and improvement of the robotaxi fleet operating on public roads. These real-world findings feed into a continuous feedback loop that leads to software improvements, additional operator training, and operational procedure adjustments, where appropriate. If an event, or a series of repeated events, generates an unanticipated safety risk above an acceptable level, Zoox may also decide to restrict, pause, or ground operations while implementing mitigations to reduce safety risk to an acceptable level. The continuous monitoring feedback loop ensures that the Zoox safety case is constantly maintained, verified, and updated with the latest information and data.[4]

Specifically related to the clearance process, operational safety centers on the tools and workflows used by our TeleGuidance team, i.e., the remote support tacticians who provide high-level assistance when the robotaxi encounters a particularly complex situation and proactively asks for assistance. TeleGuidance tacticians do not directly drive the vehicle; instead, they offer guidance such as approving or suggesting an alternate route, while the robotaxi remains fully responsible for all driving decisions.

To ensure the safety of these tools and procedures, Zoox applies structured hazard analysis such as STPA and historical learnings from operations to identify potential causes of safety risk. The resulting risk scenarios then inform the user interface design to enable tactician effectiveness. They also become testing scenarios to guide tactician training and evaluate tactician effectiveness, creating a direct link between identified risks and the readiness of the people who manage them.

The same structured approach extends to our broader operational processes. Our risk quantification methodology accounts for the possibility of both human error and tool malfunction within TeleGuidance, and these estimates feed directly into the overall CIF model.

4. For more detailed information, please see the Zoox Safety Report "Operational Safety" on our Safety page.

Together, these methodologies feed into the overall safety case and CIF metric: the safety of our driving software, the safety of our robot platform, and the safety of our operations are each quantified and rolled into one comprehensive system estimate. That estimate must meet or exceed our quantitative safety targets before any software release, hardware change, or operations update is cleared to operate on public roads.

Conclusion

At Zoox, safety is foundational. Our vertically integrated model, including robot platform, software, and operations, enables end-to-end safety ownership. Through our safety case methodology, we hold ourselves to quantitative safety targets anchored in real-world human driving data and significantly safer than the human benchmark.

As our technology matures and our operational footprint expands, we will continue to raise the safety bar, broaden and enhance our validation methodologies, and share our progress with regulators, riders, and the public. We believe transparency is essential to building the trust that the safe deployment of autonomous vehicles demands.

For more information about safety at Zoox, you can visit our safety page.