

Safety Case Framework Development and Tailoring

Development of the Framework

Aurora has adopted a safety case-based approach because we believe that it is the most logical and efficient manner to show and explain how Aurora determines that our self-driving vehicles are acceptably safe to operate on public roads. The heart of this framework is a structured argument, supported by evidence to demonstrate a claim for why our vehicles are acceptably safe. No single piece of evidence captures the totality of safety. There are complex interactions and relationships between the many elements that go into a self-driving vehicle. Ultimately, evidence without a claim is simply trivia and, conversely, a claim without evidence is baseless. A safety case-based approach brings these two essential concepts together in a logical manner to effectively show the work that we are doing to determine our vehicles are safe to operate on public roads. Along with delivering a safe product, being transparent with our approach is an important part of developing self-driving technology.

The Aurora Safety Case Framework is extensive and can be challenging to digest. This supplement provides context and insight about our process and intentions. We hope our industry, partners, and the public have a rich understanding of Aurora's thought process behind the framework and how we intend to use it as an integral part of delivering a commercial Driver.

Top-Level Claim

The Aurora Safety Case Framework is centered around the single top-level claim that ***"Our Self-Driving Vehicles are acceptably safe to operate on public roads."*** We use the entirety of the safety case to substantiate this top-level claim.

The first step in decomposing that ***Our Self-Driving Vehicles are acceptably safe to operate on public roads.*** is to address the breadth required to substantiate this top-level claim. Aurora has identified five safety principles that embody the scope of the top-level claim and provide themes to break the claim down further.

Safety Principles

Our top-level claim is broken down into the following five safety principles:

1. Proficient
2. Fail-safe
3. Continuously improving
4. Resilient
5. Trustworthy

1. Proficient

A Self-Driving Vehicle cannot be considered safe to operate on public roads unless it is suitably proficient. Proficiency includes the design, engineering, and testing necessary to develop a product. This safety principle contains the Self-Driving Vehicle performance requirements for nominal, off nominal, and corner cases of Self-Driving Vehicle behaviors.

2. Fail-Safe

The fail-safe principle addresses how the Self-Driving Vehicle behaves in the presence of faults and failures. No system is ever 100% perfect; components will wear out or have premature failures from time to time. The Aurora Driver is designed to detect and safely mitigate these failures. This safety principle contains all of the fault detection, mitigation, and notifications built into the vehicle.

3. Continuously Improving

The continuously improving principle outlines how we are enshrining the concept of continual improvement into the development of our system. A Self-Driving Vehicle is equipped with sensors, and a fleet of Self-Driving Vehicles captures significant amounts of data from just a single day's operations. We are able to harness the power of this data to enable continuous improvement. This field data feeds a comprehensive data analysis effort that calculates safety performance indicators and also considers data collected during design and development. This approach to systematically collecting and analyzing data allows us to spot trends, regressions from the mean, and emergent behaviors. Aurora also takes a proactive approach to continuous improvement, using risk identification techniques to proactively identify risks.

4. Resilient

Self-Driving Vehicles are designed to safely operate on public roads, but this does not isolate them from malicious actors or unavoidable events. The resilient principle showcases how the Aurora Driver is capable of withstanding adverse events and intentional misuse and abuse.

5. Trustworthy

Aurora's Self-Driving Vehicle may be Proficient, Fail-Safe, Continuously Improving, and Resilient,

but without the trust of the public and governmental regulators, we cannot fully realize our top level claim. The trustworthy safety principle addresses how Aurora plans to gain trust through public, government, and stakeholder engagement, safety transparency, safety culture, as well as external review and advisory activities.

Decomposition of Safety Principles

Now that our top-level claim is defined in terms of the safety principles covering the scope of safe operations, we will break down each safety principle using a breadth first, depth second approach.

Each safety principle is broken down through layers of intermediate arguments, contexts, and strategies. The lowest level claims are ultimately satisfied by evidence produced by our employees. This approach enables the tracing of each safety argument as a logical decomposition from a broad concept down to specific tangible evidence supporting the claims.

The evidence used to support a claim will come in two forms — product evidence and process evidence. Product evidence includes deliverables such as technical specifications, test plans, and test results. Process-related evidence demonstrates that the product evidence was generated in a systematic manner with sufficient rigor, review, and independence. This evidence may include informal internal audit reports confirming that we are following our established processes. Both types of evidence are needed to sufficiently address the claims in the safety case.

Across the Product Lifecycle

Aurora defines the product lifecycle as starting with our product team, and progressing through engineering, design, development, testing, manufacturing, operations, and finally ending with decommissioning. Aurora has scoped the Safety Case Framework to consider safety across the full lifecycle. This means that we are developing safety cases for our development vehicles on the roads today with vehicle operators, as well as for our future driverless commercial operations.

Across the Enterprise

The Aurora Safety Case Framework is scoped to Aurora as an entire enterprise. This means we are considering more than just the Self-Driving Vehicle in the scope of our safety cases. We are including the off-board systems, people, processes, and culture of the organization bringing the Self-Driving Vehicle to the public roads. This enterprise-wide approach ensures that we are not only considering how the self-driving vehicle itself performs, but also how we, as a company, are able to iterate on its development and refinement.

Application of the Framework

The Safety Case Framework is a tool that we use to inform the day-to-day activities of hundreds of Aurora employees on our path to develop the Aurora Driver.

The Safety Case Framework is designed to be adaptable to different vehicles, scenarios, and environments. We will use the Safety Case Framework to create a specific safety case, taking care to define its specific context and application in each instance. Think of the framework as a generic blueprint from which various specific safety cases are generated. For example, we create safety cases for specific vehicles and vehicle configurations (e.g., our truck and passenger car vehicle platforms) as well as specific operational design domains (e.g., highways). As a result, rather than a single safety case that will cover all uses of our Self-Driving Vehicles, we will have multiple individual safety cases that will cover the various configurations, platforms, and operating domains.

We also will tailor safety cases based on whether we are testing on the road with Vehicle Operators (VOs) monitoring the Aurora Driver, whether we are on a private closed track without VOs, or whether we are on public roads without VOs. Given this context, certain claims that are relevant for VO operations, such as “G1.4.1.4.2 Access to vehicles and vehicle keys are restricted to qualified vehicle operators,” do not apply in a No Vehicle Operator context. Of course, when that shift in context occurs, other claims that we may have been previously deemed non-applicable may then become germane. This tailoring of the framework is unique to each safety case being developed depending on the assumptions, claims, and arguments we make. As a result, while the Safety Case Framework may be universal, the tailoring is essential.

The Framework as a Tool

Building a commercial-ready Self-Driving Vehicle is a complicated engineering endeavor. Aurora’s Safety Case Framework is a robust and powerful tool that can be used to define and manage this complex challenge. The framework can also be used to communicate assumptions and intentions in a rational and logical manner to help the reader to understand and digest the inherent complexities. As with many tools, the outcome ultimately depends on how the user wields it. We expect that the Safety Case Framework will continue to inform our efforts and progress toward meeting the claim that we believe our self-driving vehicles are acceptably safe to operate on public roads.