Driverless Safety Report 2025

Aurora

Table of contents


AUTONOMOUS
VOLVO AUTONOMOUS SOLUTIONS
VOLVO

A letter from CEO and co-founder Chris Urmson

At Aurora, we are nearing the launch of our autonomous trucking product and, with it, are on the cusp of a new era of safer, more reliable transportation.

Since our founding, we have been driven by our mission to deliver the benefits of self-driving technology safely, quickly, and broadly. To do so, we've structured our business around safety—making it a leading factor in corporate decisionmaking, product development, fleet growth, and geographical expansion. We staunchly believe a thoughtful, pervasive, and vigilant safety approach that spans both product and organization is essential for responsible deployment of autonomous vehicles.

Now, we are publishing our 2025 Driverless Safety Report to transparently show our work and share safety expertise prior to initiating driverless operations. Throughout this report, we provide details about when, where, why, and how our product, the Aurora Driver, functions safely. We explain how we tackle risk management, redundancies, cybersecurity, remote assistance, and so much more.

Publishing detailed safety information is part of our commitment to earning the trust of all stakeholders with a vested interest in autonomous trucking. This includes government leaders and first responders, whose mandate to improve road safety aligns with our own. It includes our valued customers, who diligently improve the strength and efficiency of our supply chain every day. And, crucially, this includes the communities in which we operate, like in Texas where many consumers and businesses have already received goods hauled by Aurora Driver-powered autonomous trucks.

I am tremendously proud of the work our team has done to date, and am even more excited for what is to come. Our efforts build toward a shared vision for a safer transportation ecosystem and, in this moment, I have never been more confident in the promise of self-driving technology. Commercial launch will be a big milestone, and it is only the beginning for Aurora—our commitment to safety must be ubiquitous as we refine our product and expand to new frontiers.

As we take this next step, we invite you to join us on the journey to make our roads a safer place for everyone. We all have a stake in the future of transportation.

Chris Urmson
CEO & Co-founder

Introduction

Introduction

The National Highway Traffic Safety Administration (NHTSA) has encouraged entities engaged in developing and deploying an automated driving system (ADS) to publish a Voluntary Safety Self-Assessment (VSSA) demonstrating how they address 12 priority safety design elements.

With this fourth edition of Aurora's VSSA, we continue to show how those 12 safety elements are reflected in our overall safety approach and provide additional detail on our approach to building and operating safe autonomous vehicles.[1]

We organized this report by five major principles and in alignment with our Safety Case Framework. This organization allows us to account for the development of our autonomous vehicle as a product, and also discuss the process and company-wide organizational factors that support each principle.

01 Ensure Proficiency

This principle focuses on our autonomous vehicles' capabilities, and the software, sensors, and computing systems necessary to enable autonomous driving under nominal conditions. In this section, we also discuss how we leverage our extensive virtual testing and monitor on-road performance to improve the Aurora Driver.

02 Fail Safely

This principle details how the Aurora Driver will respond when something goes wrong. This section emphasizes our Fault Management System, which is responsible for ensuring that the vehicle will be able to take appropriate action if a failure should occur.

03 Continuously Improve

This principle and section outline that we have the appropriate safety metrics and safety performance indicators in place to provide feedback and early warning to enable us to proactively address issues before they can cascade further.

04 Operate with Resilience

This principle focuses on how we have designed the Aurora Driver to withstand or recover from threats. In this section, you will find our cybersecurity approach and how the vehicle will respond in the event of an on-road incident.

05 Be Trustworthy

This principle focuses on ensuring that Aurora is a responsible company that embodies safety. We discuss how our work as a company supports this principle in this section.

[1] References to Aurora within this document mean Aurora Operations, Inc. and its affiliated entities.

Executing our safety case

At Aurora, we are using a safety case-based approach to inform, guide, and determine that our technology is acceptably safe to operate on public roads.

A safety case is a logical argument, supported by evidence, intended to justify that a system is acceptably safe for a specific application in a specific operating environment. A structured argument includes a specific claim—such as that our self-driving vehicles are acceptably safe to operate on public roads—that is broken down into sub-claims, which are then ultimately supported by evidence. We believe a safety case is the most effective and efficient path to safe driverless operations, and is imperative for any company looking to safely deliver commercial-ready self-driving vehicles at scale.

Safety case-based approaches have been widely used in other safety-critical industries like nuclear, aviation, rail, and medical devices. The term "safety case" is also used in some existing voluntary automotive industry standards. We have adapted our safety case approach based on the best practices of these industries, and applied it to developing and operating autonomous vehicles. We have developed a Safety Case Framework that includes the claims we believe are necessary to assert that our vehicles are acceptably safe to operate on public roads across various use cases and Operational Design Domains (ODDs).[2]

From this overarching framework, we are developing and maintaining specific *tailored* safety cases that each contain the claims relevant for a specific use case (e.g., autonomous driving with vehicle operators or autonomous driving without vehicle operators within a specific ODD).

For example, when developing the Aurora Driver, we currently have vehicle operators in our vehicles monitoring the performance of the Aurora Driver at all times and ready to take over as necessary to ensure operational safety. Therefore our tailored safety case for this use case includes claims focused on vehicle controllability and vehicle operator hiring, training, and operational procedures, among others. However, as we approach the point of removing vehicle operators from the vehicles, these vehicle operator-centric claims will no longer be relevant. At that point, we will have completed a tailored safety case to include other claims from the Safety Case Framework related to demonstrating acceptably safe driverless operations within our ODD.

Currently, we have completed several safety cases for different vehicle platforms that require vehicle operators to be in the vehicle to supervise the Aurora Driver. With this report update, we are on the verge of closing our driverless safety case for our initial launch route between Dallas and Houston, as described later in this report. We will continue to maintain multiple safety cases, each tailored for different use cases as appropriate and building on prior safety cases that have already been closed.

The Aurora Safety Case Framework takes into account the entire development lifecycle of our vehicles—from testing and operations involving vehicle operators to driverless operations. We have adopted a safety case-based approach because we believe it is the most logical and efficient manner to explain and demonstrate that our self-driving vehicles are acceptably safe to operate on public roads.

[2] https://aurora.tech/blog/aurora-unveils-first-ever-safety-case-framework.
See the Ensure Proficiency section for an explanation of ODD.

Aurora Driver for Freight

Aurora's core product is the Aurora Driver. The Aurora Driver consists of the hardware, software, and data services required to safely and productively deploy and operate self-driving vehicles.

Aurora is building the Aurora Driver to scale across diverse vehicle platforms, and is prioritizing heavy-duty Class 8 trucks that haul trailers as our launch product. We call this product Aurora Driver for Freight, which combines the Aurora Driver with a set of services that provide carriers and private fleets with a reliable and scalable driver supply seamlessly integrated into their existing operations. The Aurora Driver is currently capable of autonomously hauling different trailer types; for example, we regularly haul dry van and refrigerated trailers, along with intermodal trailers.

In the coming year, the Aurora Driver for Freight will operate between terminals situated outside of dense urban areas, which enables us to aggregate loads from multiple customers into one place at either end of a major freight route. This approach reduces the extent of our off-highway driving and allows Aurora to control launch and landing operations, which includes pre- and post-trip inspections, service, and maintenance. Aurora is exploring delivering loads directly to select customers' endpoints in the next year as well.

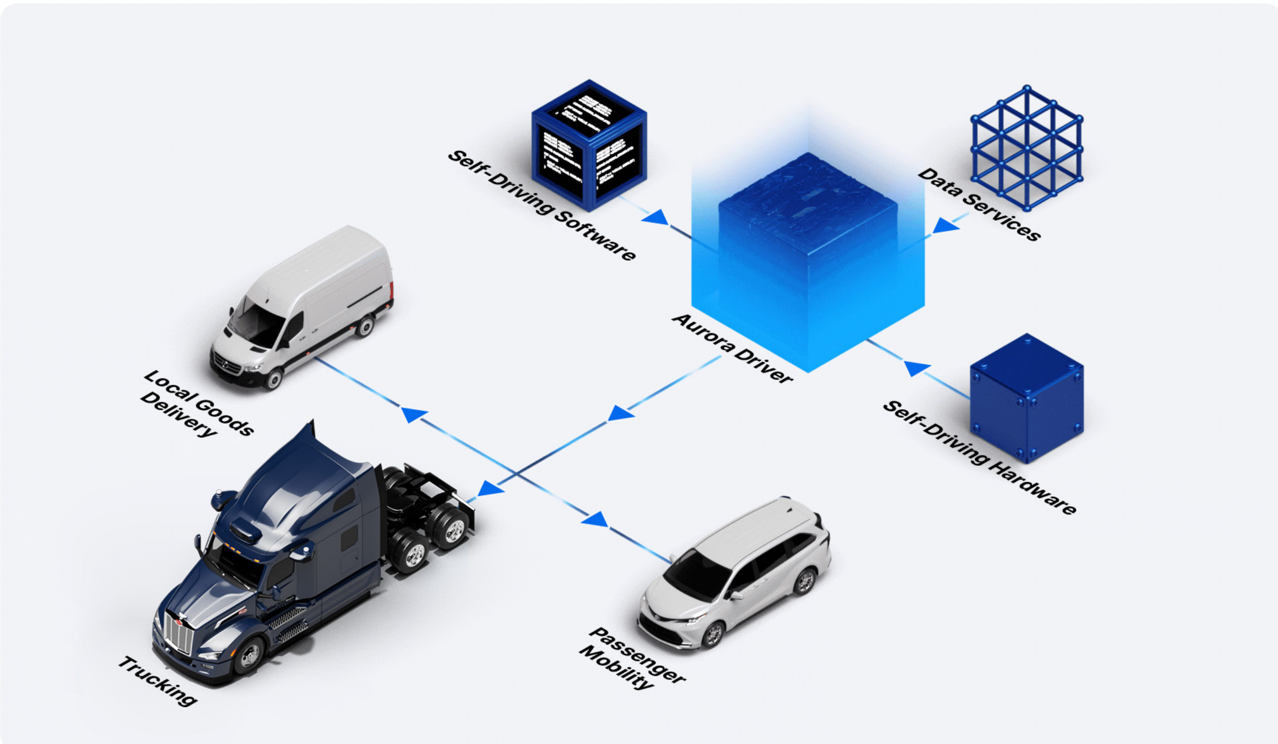

01 | Ensure Proficiency

1.1

The Aurora Driver

The Aurora Driver consists of the hardware, software, and data services required to safely deploy a self-driving vehicle, and we are designing the Aurora Driver so that it can operate with diverse vehicle platforms. This unified technology architecture not only allows us to leverage our development work across multiple use cases, but also enables us to scale across various vehicle platforms while maintaining consistent capabilities and performance.

1.1.1

Perception

Driving safely and proficiently requires having accurate, detailed, and timely information about the surrounding environment.

High-quality perception sensor data is the foundation for the Aurora Driver. Our state-of-the art perception system leverages deep learning and sensor fusion to see the world in full 3D and at long ranges. This system provides the Aurora Driver with situational awareness, made possible through our innovative approach to the integration of advanced machine learning and probabilistic state estimation. The Aurora Driver uses a combination of camera, radar, and lidar sensors to maximize coverage and take advantage of each sensor's unique capabilities.

1.1.1.1

Cameras

Cameras observe visible light, and we equip the Aurora Driver with cameras distributed around the vehicle. Our cameras' high resolution sensors and quality optics enable the Aurora Driver to operate both on high-speed highways and in congested urban settings. Our camera system enables the Aurora Driver to "see" the world around it reliably and in great detail, allowing it to recognize the color of a stoplight, interpret traffic signs, discern lane markers, and spot other actors in the environment. We have also incorporated high dynamic range capabilities to compensate for differences in brightness needed when driving in quickly-changing light conditions, like exiting a tunnel or passing through the dark shadow of an overpass on a sunny day. We custom designed the lensing, layout, and cleaning solution for our cameras to meet the demands of the broad set of use cases in which the Aurora Driver will operate.

1.1.1.2

Radar

Radars emit radio waves and measure the return of those emitted waves reflected from objects in the scene. We use cutting-edge imaging radars to provide both elevation and azimuth data, helping the Aurora Driver quickly and accurately perceive the objects around it. Radar is particularly robust to adverse weather conditions that challenge camera and lidar, like rain and fog. When combined with the other sensors, the radars we use provide overlapping sensor inputs to the Aurora Driver.

1.1.1.3

Lidar

Lidar systems allow self-driving systems to see and understand their surrounding environment. Aurora uses two types of lidar—conventional Amplitude Modulation (AM) lidar and our proprietary FirstLight lidar. While cameras supply high-resolution 2D images, lidars provide texture by adding depth to the scene, effectively allowing the vehicle to sense the world more accurately, in 3D.

Conventional AM lidar works by emitting brief light pulses at a fixed frequency. This helps determine the location of objects at short and mid range based on how long it takes for those laser pulses to bounce off surfaces and return to the sensor—the farther away something is, the longer it takes for the light to return.

Aurora's industry-leading, proprietary Frequency Modulated Continuous Wave (FMCW) lidar, named FirstLight, is able to detect and track objects at long range—more than 450 meters away. FirstLight lidar allows the Aurora Driver to perceive an object nearly nine seconds earlier than conventional lidar for vehicles moving at highway speeds. This gives the Aurora Driver a huge advantage—because it can sense farther, it has more time to react to unexpected obstacles. FirstLight lidar also experiences less interference from the sun or other lidars, because each sensor is designed to respond to only the signals it creates, making the data received more relevant and less noisy.

FirstLight's technology also allows us to measure the relative velocity of other objects instantaneously. It can distinguish when another vehicle is pulling away from us or if we are closing in on it, providing data to the perception system so it can generate a more reliable and rapid estimate of velocity and enhancing the Aurora Driver's ability to distinguish between actors that are moving at different speeds. Faster object recognition and tracking by the perception system allow the motion planning system to react better and give the Aurora Driver more time to maneuver. With FirstLight, the Aurora Driver can detect and track objects more quickly and at great distances, ultimately making quicker decisions with greater accuracy—a significant advantage while driving at highway speeds.

1.1.1.4

Degraded sensor monitoring & cleaning

The Aurora Driver's perception system is constantly assessing the range and quality of the signal its sensors observe. In the event of sensor fouling (due to environmental conditions such as rain or dust), the Aurora Driver is capable of detecting sensor degradation, cleaning the sensor, and ultimately pulling over if cleaning proves to be insufficient. We have designed an innovative cleaning system that prevents moisture and debris from building up on our sensors' lenses. Each of our sensors is equipped with special nozzles that blast a combination of high-pressure air and washer fluid across their surface—cleaning each lens in a fraction of a second. If conditions are degraded too much for normal driving, the Aurora Driver will slow down. In severe cases in which conditions fall outside the Aurora Driver's ODD, the Aurora Driver will take action to achieve a minimal risk condition, which could include finding a safe location to pull over and wait out the storm.

1.1.2

Planning

Driving is full of complex interactions. How one driver behaves on the road affects how other vehicles and road actors in the environment behave.

The Aurora Driver is designed to act safely around others on the road and to be predictable. We fuse machine learning with formal rules to create a robust motion planner that allows the vehicle to smoothly navigate situations while retaining the ability to operate in a safe and predictable manner.

On the road, there is a clear set of rules (e.g., that the Aurora Driver must stop at red lights) that we encode as “invariants”, and between these invariants, our system learns how to drive well based on observing how our expert drivers behave. By combining the best of modern AI approaches with invariants, we are able to build a driver that is both predictable in its behavior and trained to follow the rules of the road. This is what we call “Verifiable AI”. We have explained in depth on our blog why we are taking this approach, how we prioritize AI transparency, how we prioritize AI alignment, and how this benefits the scalability of our product.

1.1.3

Vehicle control

Once the motion plan has been generated, the Aurora Driver executes the planned actions.

This includes adjusting the vehicle's steering, acceleration, and braking to maintain the planned trajectory. Driving is a dynamic activity, and the Aurora Driver is constantly observing the surrounding environment through the perception system and making updates to the motion plan in response to others on the road, whether other vehicles speed up/slow down, enter/exit our lane or adjacent lanes, or whatever else may be encountered on the roadway. Vehicle control responds to the plan as it changes, based on new information.

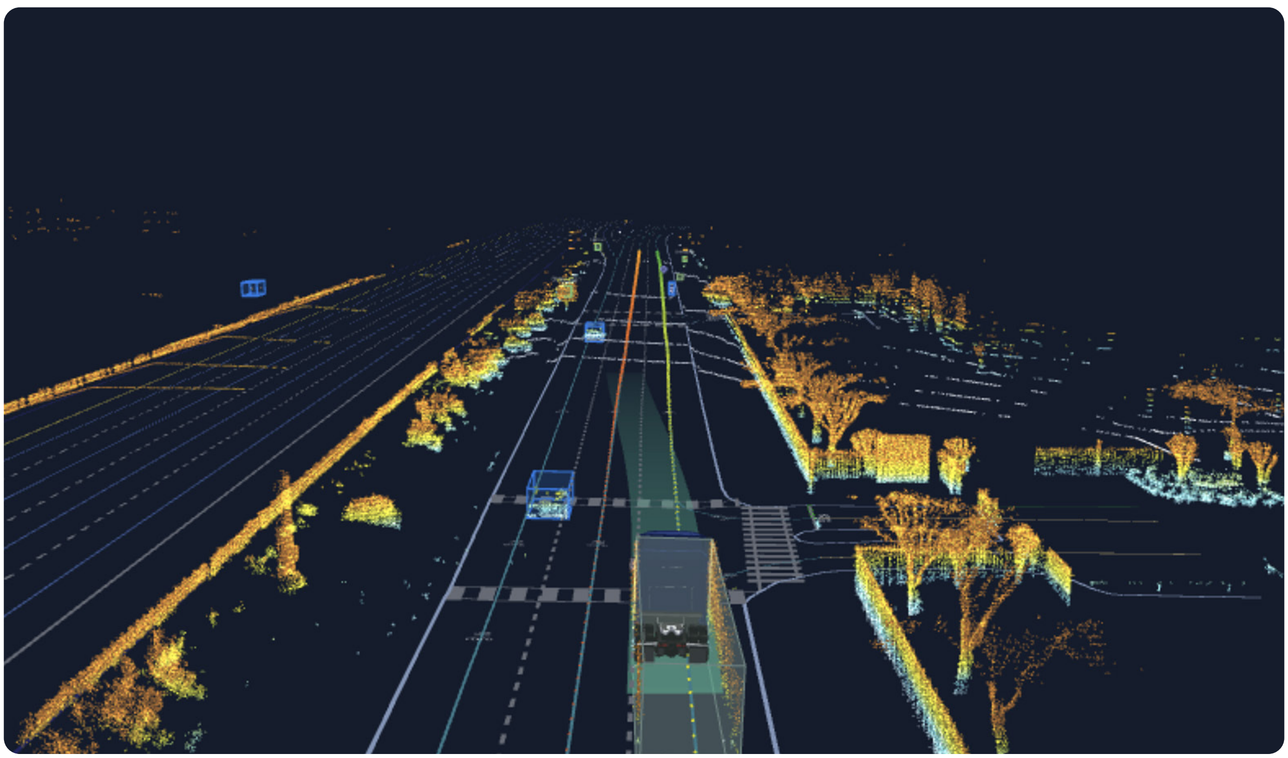

1.1.4

Mapping and localization

An accurate map helps the vehicle in many ways—it allows the Aurora Driver to locate itself precisely in the world and understand the locations of fixed objects like traffic lights, speed limits, one-way streets, and traffic circles.

The Aurora Atlas is our proprietary map for the Aurora Driver and contains more detailed information than the maps used in typical in-car navigation systems. We designed the Aurora Atlas to not only allow us to place map content very accurately relative to the self-driving vehicle, but also to allow the map to be updated frequently, quickly, and efficiently. When vehicles encounter unexpected road features, like a closed lane, we are able to update the Aurora Atlas and share that information with all other vehicles operated by the Aurora Driver.

Our onboard localization system estimates the state of the vehicle by fusing data from multiple sensors and other software modules. By considering multiple sources of information, the system can maintain an accurate estimate even in the presence of challenging factors like sensor failures, heavy traffic, challenging weather, and construction.

1.2

Aurora's Operational Design Domain

The Operational Design Domain, or ODD, refers to the conditions in which the Aurora Driver is designed to function—including the geography, environment, weather, driving conditions, and other considerations. The ODD is a critical aspect of defining capabilities because it outlines the limits of where and in what conditions the system is designed to operate.

In designing the Aurora Driver, we start by looking at the domain in which we want to operate. We then define the ODD based on the roadways, conditions, use cases, restrictions, and scenarios within that domain. From there, we begin simulations of scenarios we expect to encounter within this ODD. Our trained vehicle operators drive the roads manually to collect data and gain firsthand understanding of the ODD. Our engineers then review the collected data to develop the autonomy capabilities needed to operate in the ODD. As we continue development and testing, the engineering team continues to refine and make any necessary software modifications needed for autonomous operations within the ODD.

When deployed as a commercial product, the Aurora Driver will only operate in areas in which its proficiency has been validated. At the software level, we impose a system prohibition on operating autonomously in areas outside of the ODD, including restricting routing of the Aurora Driver based on a set of configurable elements such as road speed, road type, and traffic-control devices.

We also use a variety of triggers to detect and avoid out-of-ODD conditions for which a restriction may not have been created. While the Aurora Driver is able to detect and autonomously manage certain conditions on its own, these situations may also include review by a remote assistance specialist in response to a vehicle trigger, which will often correspond with a vehicle action to achieve a minimal risk condition (see section 4.3 about our Command Center and approach to remote assistance).

This combination of constraints in any given ODD serves as an important control to manage safety risks at every stage of development, testing, and operations.

At launch and as we progress our technology in 2025,[3] we expect our vehicles will operate in autonomy within the following ODD:

- On public roads in Texas, New Mexico, and Arizona
- Within the speed limits of the roadways where we operate (speed limits ranging from 25 to 75 mph)
- In suburban and urban areas, including in dense traffic
- Day and night
- In highway construction and work zones with cones and barriers
- In weather conditions for which we understand how the Aurora Driver will perform, including light to medium rain

[3] For additional information, please see Aurora's roadmap

1.2.1

Roadway types

The Aurora Driver is capable of driving on asphalt, concrete, and gravel. We are able to navigate all road types and features on our route including:

A variety of roads, such as

- One-way and two-way roads, both divided and undivided
- Arterials, collectors, connecters, frontage, limited access, and highways

Roadside features, such as

- Bollards, curbs, hydrants, guardrails, jersey barriers, medians, gores, bridges, and over and underpasses

1.2.2

Geographic area: DAL ↔ HOU

The Aurora Driver for Freight launch lane connects Aurora's south Dallas terminal with its north Houston terminal.

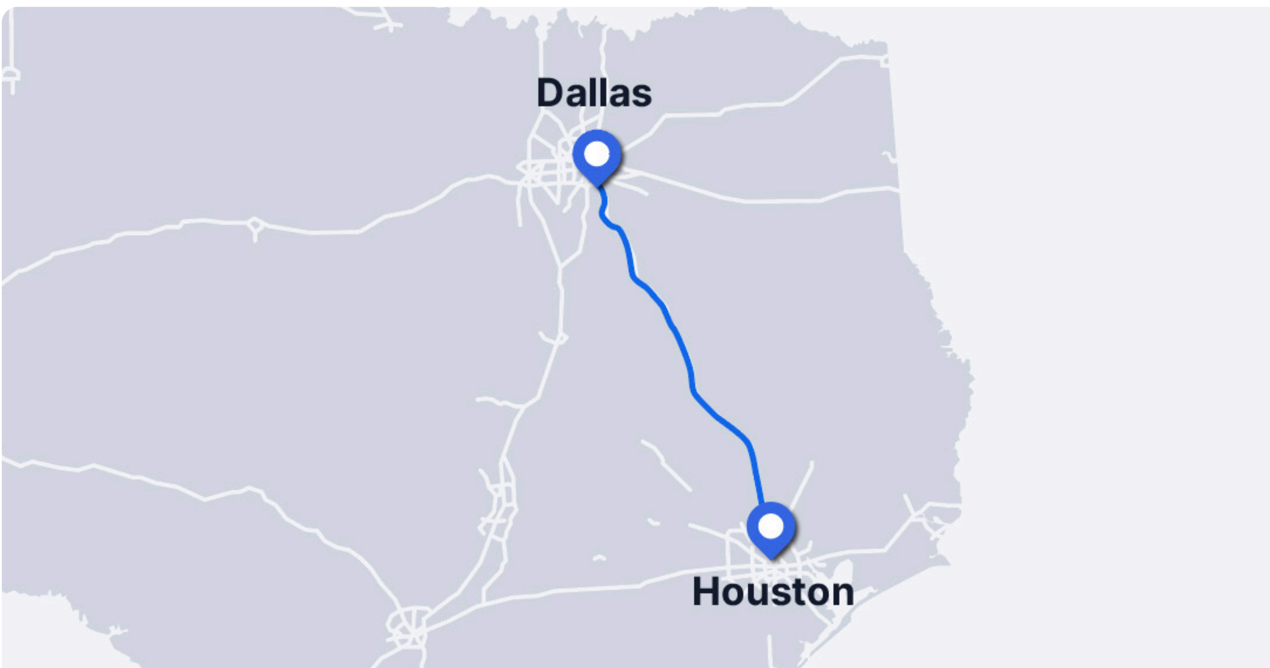

Graphic. DAL ↔ HOU endpoints

1.2.3

Geographic area: FTW ↔ ELP

Aurora also autonomously hauls commercial freight between terminals located in Fort Worth and El Paso, Texas.

Graphic. FTW ↔ ELP endpoints

1.2.4

Geographic area: ELP ↔ PHX

Aurora additionally operates autonomously between El Paso, Texas and Phoenix, Arizona.

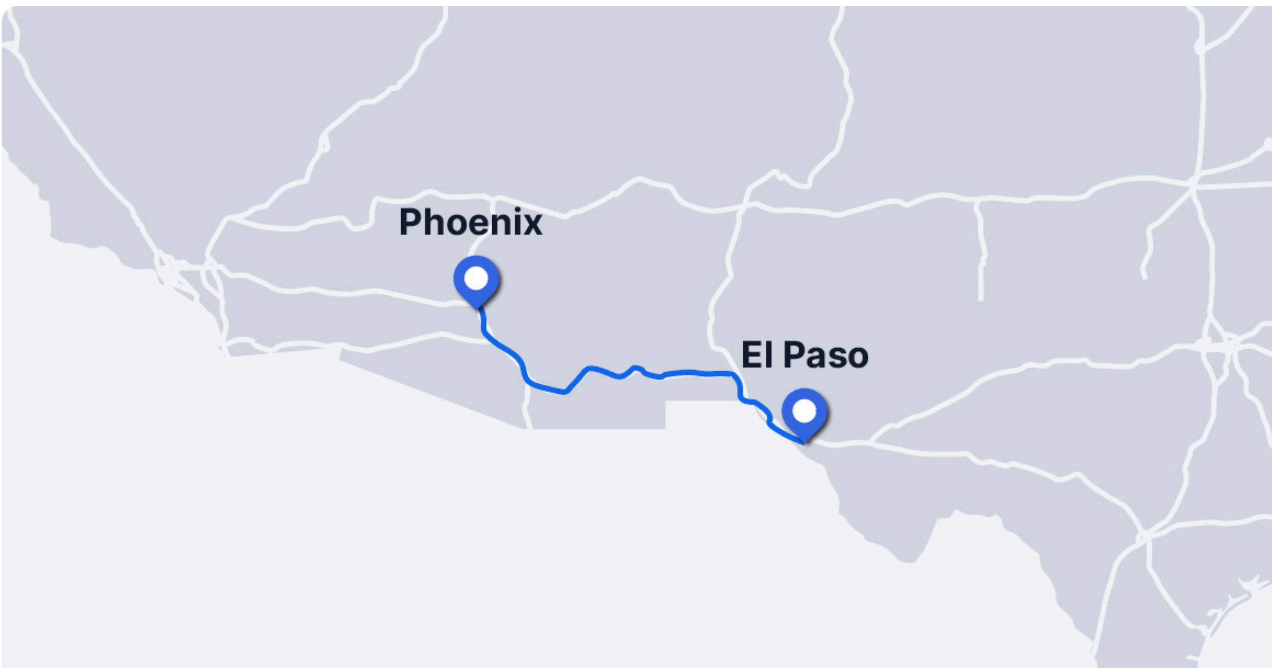

Graphic. ELP ↔ PHX endpoints

1.2.5

Speed range

The Aurora Driver's speed range is from 0 to 75 mph. To help realize greater fuel efficiency and safety benefits, the Aurora Driver cruises at 65 mph. When higher speeds are necessary (as when merging or overtaking), it will drive at up to 75 mph.

1.2.6

Environmental conditions

The Aurora Driver will operate up to 24 hours a day and seven days a week. At commercial launch, the Aurora Driver will be able to driverlessly operate in sustained winds up to 25 mph and wind gusts up to 35 mph. We plan to expand our ODD in 2025 to include driverless operations in inclement weather as well.

Under certain environmental conditions where the Aurora Driver's visibility range or safety performance would be degraded, the Aurora Driver will detect the off-nominal condition and respond by slowing down or pulling to a preferred exit or the shoulder, and trigger a condition that would prompt feedback from a remote assistance specialist. (Please see 1.1.1.4 Degraded Sensor Monitoring & Cleaning for additional information.)

1.3

Safety engineering

As part of the ODD characterization process, Aurora conducts a hazard analysis process. Hazards are potential sources of harm observable at the vehicle level.

The hazard identification process considers root causes related to design insufficiency, performance insufficiency, potential misuse, and faults and failures of the system. The hazard analysis directly informs the development of safety requirements, which closely analyze the system components, operations, and interactions with the world with the objective to prevent hazards or mitigate risks as a result of the hazards. Safety requirements specify safety goals, operational constraints, and safety measures and mechanisms which are verifiable. Finally, these autonomy capabilities undergo verification and validation to ensure that they meet the specifications outlined in the initial requirements.

In addition to the processes above, Aurora has established behavioral safety processes drawing from Safety of the Intended Function (SOTIF), as defined by ISO 21448, and functional safety, as defined by ISO 26262. As described in greater detail in the following sections, together these two voluntary industry standards present a method for addressing situations that exceed the performance limitations of the system, or failures of the system function, respectively, by the identification of triggering conditions for hazard mitigation. Aurora has internally assessed its conformance to a number of voluntary industry safety standards and best practices from organizations such as SAE International, the Automated Vehicle Safety Consortium (AVSC), the Institute of Electrical and Electronics Engineers (IEEE), and the Commercial Vehicle Safety Alliance (CVSA).

1.3.1

BeSafe and SOTIF

Aurora performs a failure modes-driven approach to reduce the risk of hazards in its vehicles, similar to SOTIF, which is encapsulated in our Behavior Safety (BeSafe) Analysis Framework.

The outcome of the BeSafe analysis is a comprehensive accounting of the triggering conditions which cause hazardous behavior that exceed the performance limitations of the Aurora Driver, and how they are monitored and mitigated when necessary. The figure below represents the initial condition, before we complete the BeSafe analysis.

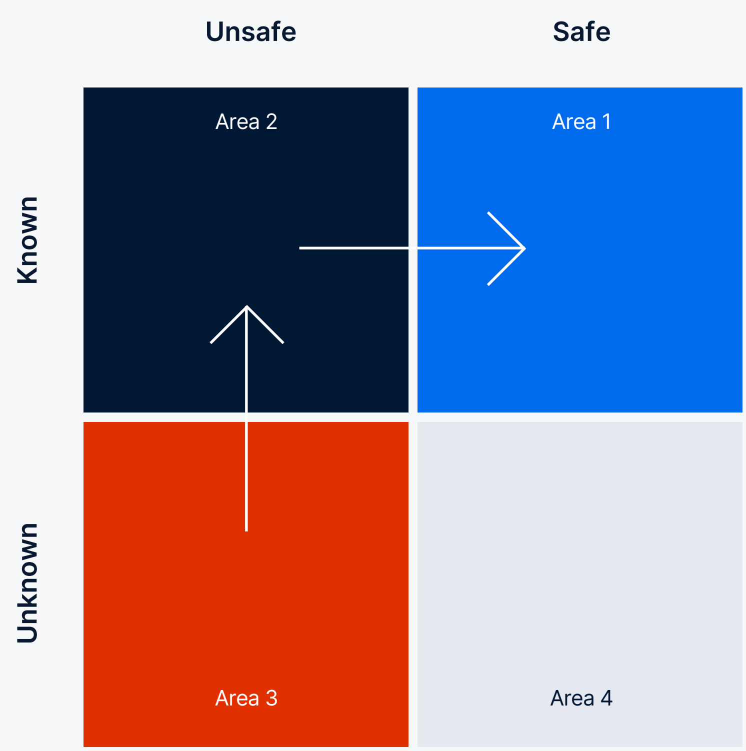

Figure. Scenario categories and the desired evolution derived from SOTIF analysis

Area 1

represents the known-safe scenarios. These are scenarios that the ADS is capable of handling safely. These scenarios may fall well within the ODD or may represent scenarios that have already been addressed through existing countermeasures. For example, this may involve the Aurora Driver maneuvering around a disabled vehicle on the shoulder, which is a scenario we encounter on a daily basis and the Aurora Driver is designed to guard against.

Area 2

represents known-unsafe scenarios. These are scenarios known to the ADS developers in which the ADS cannot perform safely. Triggering conditions relevant to this area may be identified through analysis and guided brainstorming, field data, and crash analysis and reconstruction. The objective for these scenarios is to identify countermeasures, such as functional or architectural changes to the ADS or restriction of use cases to reduce the overall risk. As these countermeasures are integrated into the system design and demonstrated to effectively mitigate the known-unsafe scenarios, these mitigated scenarios from Area 2 can be reclassified under Area 1 as safe. For example, this may involve operating in icy conditions because we know that icy conditions can occur and the Aurora Driver is not currently designed to operate in them. As a result, we operationally control our exposure to them by not deploying the Aurora Driver when roads are icy.

Area 3

represents unknown-unsafe scenarios. These are scenarios unknown to the ADS designers and under which the ADS cannot perform safely. The unknown-unsafe events are discovered through a combination of approaches, such as analysis and guided brainstorming, on-road and track testing, and simulation. As more unknown-unsafe scenarios are identified, the mitigated scenarios are then reclassified under Area 1 or 2. Area 3 is the core challenge in SOTIF, and the scenarios within contribute to the residual risk. It is inherently difficult to conjure an unknown-unsafe scenario, but an example might be if a developer had never considered that a truck could haul oversized loads like windmill blades. If this were never considered, and the autonomous vehicle did not track these vehicles properly, it may end up with unsafe behavior. That said, for reference, the Aurora Driver has encountered this scenario previously and knows how to respond to such oversized loads.

Area 4

represents unknown-safe scenarios. These are scenarios that the ADS can safely handle through existing countermeasures, even though they are unknown to the designers. An example of this may be a helicopter landing on the highway—while the Aurora Driver was not explicitly trained to detect aircraft, we anticipate the Aurora Driver would be able to use its perception generics capability to track the object and maneuver around the helicopter safely.

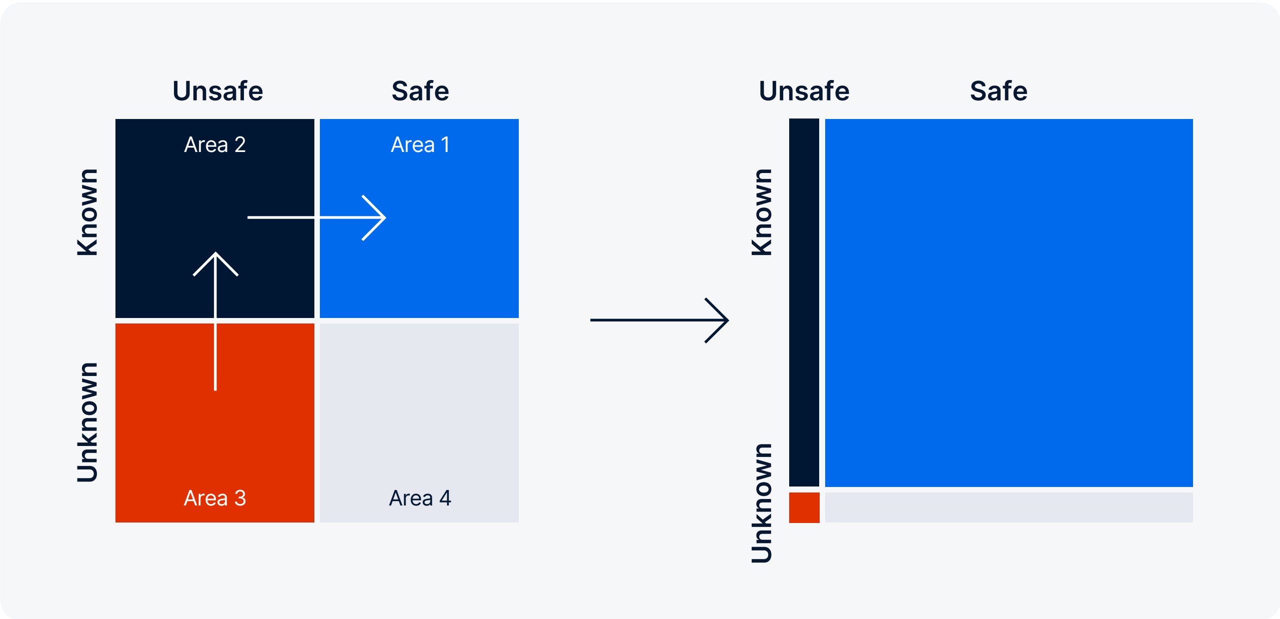

Figure. Outcome of the BeSafe analysis

As shown in the figure above, Aurora's implementation of SOTIF through our BeSafe approach enables our engineers to systematically and iteratively explore the operational space, beyond what was just encountered on road operations, to decrease the unknown-unsafe (Area 3) scenarios by shifting them into the known-unsafe space, and develop countermeasures to convert known-unsafe (Area 2) to known-safe scenarios (Area 1).

1.3.2

Functional safety

Aurora has defined a functional safety process that covers systematic, electric and electronic, mechanical, applicable environmental, calibration, and safety mechanism failures that can lead to a hazard. Starting with the identification of vehicle level hazards, these failures are defined as potential sources of harm caused by insufficient performance, misuse, design, or malfunctions, that are observable only at the vehicle level. The Hazard Analysis and Risk Assessment (HARA) helps to determine the baseline risks of a given vehicle level function within the intended use cases in the presence of malfunctions. Through this process, we are able to derive and enumerate safety goals and requirements.

1.3.3

Behavioral competencies

Based on the requirements identified for the ODD, Aurora has distilled these into a number of core autonomy capabilities. The Aurora Driver should be able to execute all of these capabilities, as necessary, in all nominal conditions. To that end, the Aurora Driver can detect and respond to the environment or actors on the highway through nominal driving interactions through the following high-level behavioral competencies:

- Adaptive cruise control
- Lane keep
- Merge
- Lane change
- Nudge
- Construction
- Intersections
- Emergency vehicles
- Vulnerable road users (VRUs)
- Debris
- Collision avoidance and response

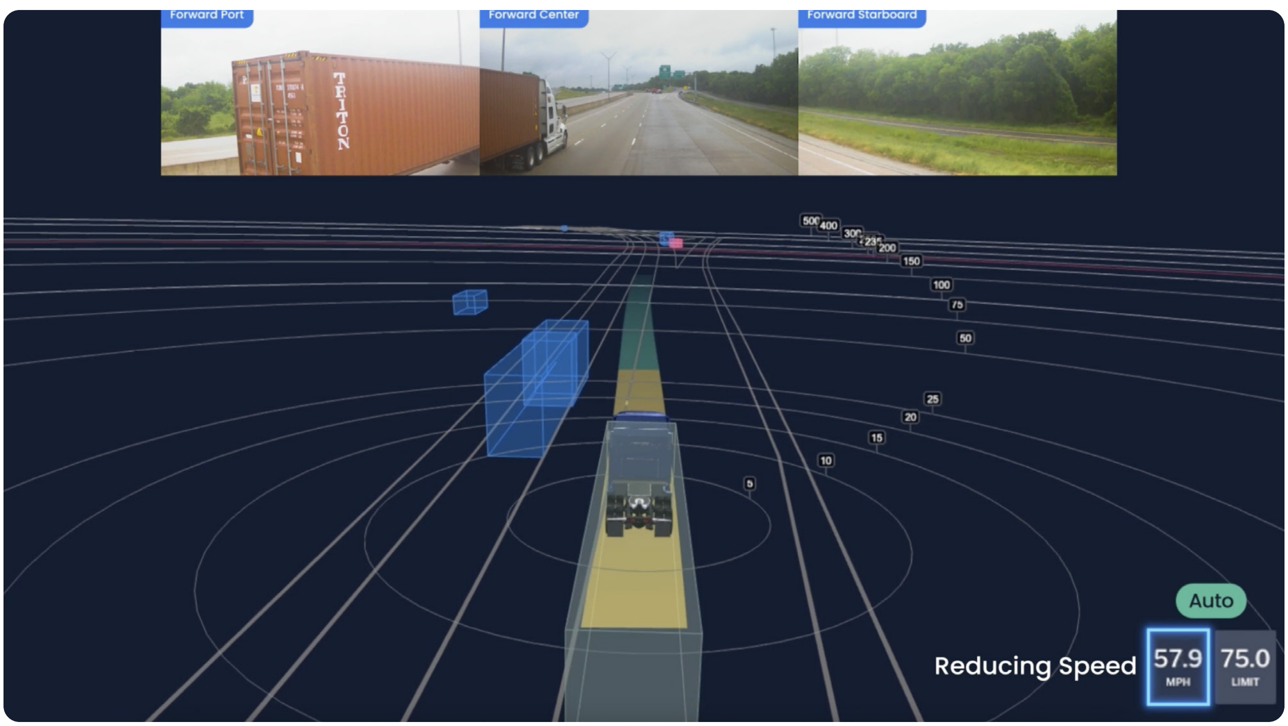

1.4

Testing the Aurora Driver

We take a rigorous approach to testing the Aurora Driver. Our virtual testing, combined with closed track and on-road testing, enables us to test the Aurora Driver in a comprehensive manner.

When validating a self-driving system, one of the challenges is that many of the most dangerous scenarios on the road are also the rarest and the least likely to be encountered. This means that relying exclusively on on-road testing is not enough. Validating safety requires a multi-modal approach—a combination of road and virtual testing, which is why Aurora's offline simulation tools are so critical.

1.4.1

Virtual testing

Though our vehicles drive thousands of miles each month, we cannot possibly test every situation that we may encounter on a real roadway—nor would it be safe to do so. So, we test the Aurora Driver virtually, by running it in simulated environments based on synthetic or real historical data. Extensive virtual testing allows us to understand how the Aurora Driver performs in millions of scenarios, far beyond what we would ever be able to experience on a test track or the road.

Aurora's Virtual Testing Suite is an essential part of our validation strategy. It enables us to quickly turn on-road events into multiple opportunities for us to learn and advance, including allowing us to both simulate and replay the data our vehicles see on the road. This means that we can create the equivalent of thousands of vehicles operating on the road at once, enabling development at a scale that would be otherwise unachievable through on-road testing alone.

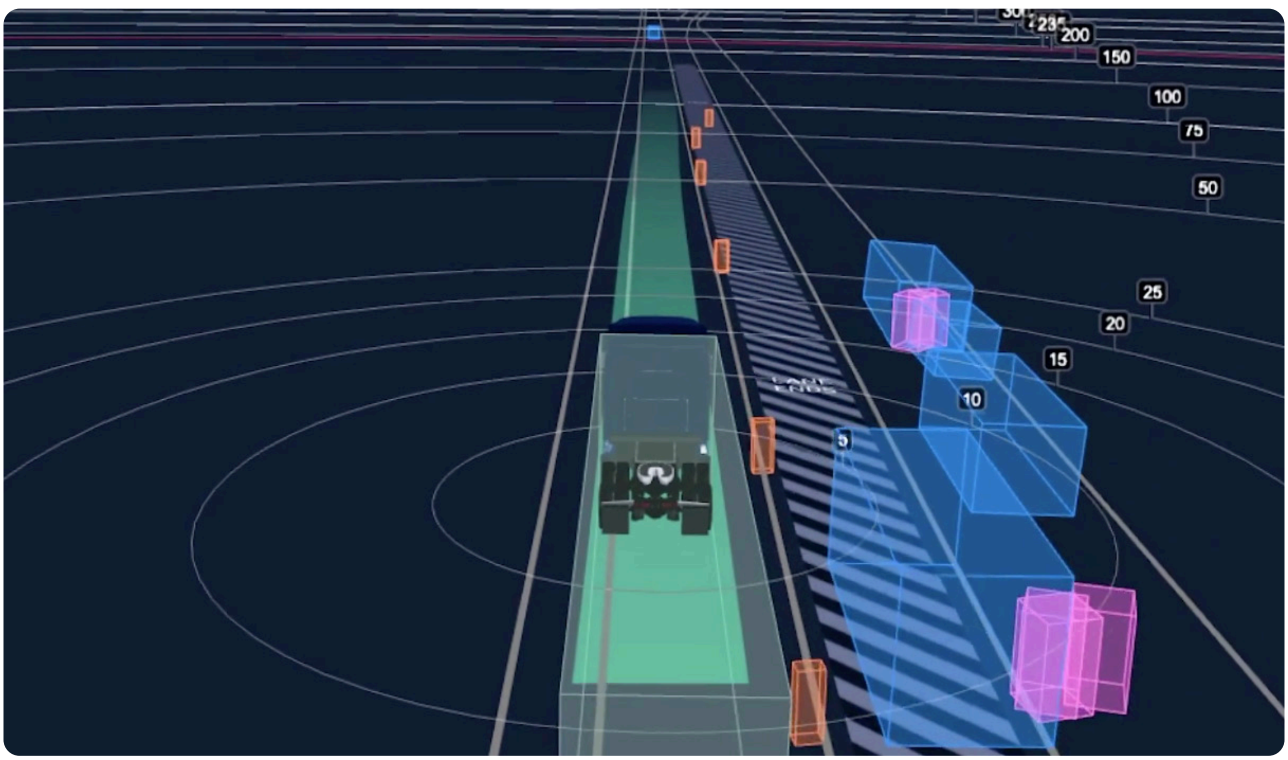

Scenarios encountered in trucking operations

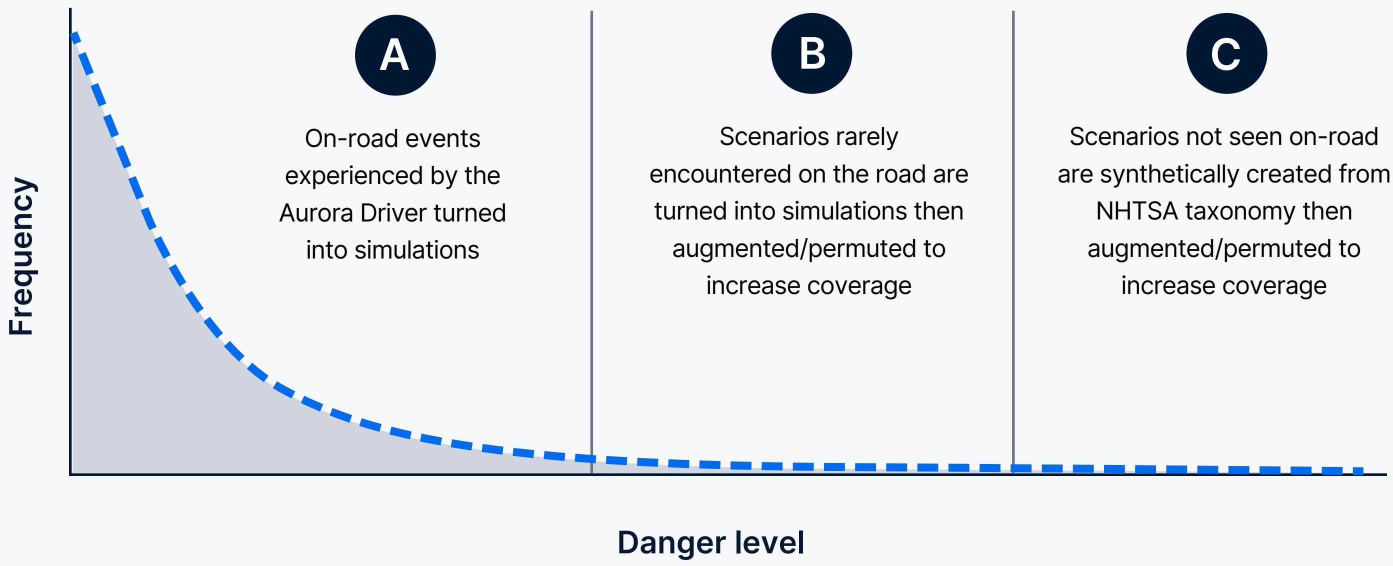

To help understand how we use the Aurora Virtual Testing Suite to amplify exposure to rare events, we break down the scenarios encountered during operation into three buckets based on how frequently they occur on the road:

- **A** These are the most common scenarios and on-road events that we see regularly, which are often turned into simulations.
- **B** These are scenarios we rarely encounter, which are turned into simulations and then augmented and modified to help ensure they are appropriately included in testing.
- **C** These are the scenarios that we have not experienced on the road, but we know that they can happen. These scenarios could also include previously unknown events described in section 1.3.1 BeSafe and SOTIF, which are events the Aurora Driver has not encountered but are identified as credible. In order to evaluate the Aurora Driver's performance in these scenarios, we synthetically generate the rarest of events based on NHTSA's imminent collision crash set, as detailed in section 1.4.1.2.

This framework provides a methodology by which we are able to build test coverage of the scenarios that the Aurora Driver will encounter on the road.

1.4.1.1

Types of virtual tests

Our approach to virtual testing uses a suite of tools that allow us to assess how software works at every level. We convert on-road events into one or more of the following types of virtual tests:

Perception tests

A perception test considers whether the perception subsystem's (described in section 1.1.1) overall performance meets the requirements for the given ODD. Ultimately, the perception subsystem needs to provide a set of capabilities that enable the Aurora Driver to operate safely within the ODD. While we continue to develop highly realistic sensor simulations to generate tests for uncommon and high-risk scenarios, most of our perception tests come from real-world information, from which we create a series of specialized perception tests.

For example, let's say the vehicle encounters road debris in the lane of travel. Specialists review footage and then label aspects like object category (e.g., road debris) and velocity (e.g., stationary or moving). We can then use that real-world point of reference to virtually test how well new versions of the perception software are able to see and identify such debris on the road.

Road-based tests

We want the Aurora Driver's movements to be natural and predictable for other drivers and pedestrians in the world. To accomplish that, we test whether our motion planning software can emulate what a trained, expert driver would do in complex situations. Our goal is to help the Aurora Driver choose a safe response to a scenario.

For example, we may want to understand when the Aurora Driver would start making a right turn compared to when our expert vehicle operators begin turning. To do this, we collect data while our vehicle operators drive and then assess how the Aurora Driver's plan compares with the vehicle operators' actions. The system is able to learn from these experiences; the ability to test future versions of our software against these expert demonstrations underscores the tremendous value in curating the right set of manual data to keep improving how we drive autonomously.

Synthetic tests

Synthetic simulations are virtual models of the world where we can test how the Aurora Driver reacts in many permutations of the same situation. These synthetic simulations also allow us to simulate a wide variety of interactions between the Aurora Driver and other actors in virtual worlds, including those that are only encountered incredibly rarely in the real world.

1.4.1.2

Imminent collision test set

As part of its ongoing efforts to understand and improve roadway safety, NHTSA has developed a detailed taxonomy of crash types that reflect all law enforcement-reported collision scenarios that human-driven motor vehicles experience on the road.[4]

Aurora has incorporated this taxonomy as a core element of our validation process—considering vehicle movement, actor placement, road infrastructure, and more—to develop the Aurora Driver so that it can react appropriately even in the most challenging of scenarios, and, in many situations, prevent collisions.

We are training the Aurora Driver to handle thousands of permutations guided by the NHTSA collision taxonomy. We build these scenarios to be extremely challenging to help ensure that the Aurora Driver performs reasonable maneuvers in all manner of situations. This is fundamental to understanding the expected safe behavior by an ADS in potential collision situations.

Many drivers may only encounter a roadway collision a few times in their life, but the Aurora Driver has experienced millions of these scenarios in simulation. These simulations have advanced our technology and built confidence in our ability to operate in the real world. Planning for how our autonomous vehicles should perform in these situations is an essential part of how we complete our safety case for driverless operations on public roads.

While NHTSA's crash taxonomy helps us anticipate what *could* happen on the road and prepare accordingly, it is also essential to understand what *has happened* on the road—providing the ability to learn from past collisions by human drivers and help prevent them in the future. As a second source of crash scenarios, we evaluated real examples of fatal collisions between human drivers on our Dallas-Houston launch route.

As part of our validation of the Aurora Driver's performance at commercial launch, we examined fatal collisions involving human-driven Class 8 trucks on our Dallas-Houston launch route between 2018-2022. Using data available about these crashes, we recreated the collisions in simulation to understand how the Aurora Driver would have acted in comparable circumstances. In the 29 instances in which the Aurora Driver could have been operating the initiating vehicle, the combination of its powerful sensors, attentive driving, and quick decisionmaking would have prevented the collisions.

[4] NHTSA Pre-Crash Scenario Typology for Crash Avoidance Research (DOT HS 810 767)

1.4.3

Track testing

We drive the vast majority of our miles in our Virtual Testing Suite, enabling us to make rapid progress on the Aurora Driver's capabilities. But high-quality real-world experience is still vital to develop and refine the Aurora Driver.

We continually test and validate our progress safely in the physical world, including on closed-course tracks that provide a controlled environment to both evaluate the Aurora Driver's performance and ground our simulated scenarios with real-world data. Track testing allows us to create specific, repeatable tests that would be near impossible to encounter or potentially dangerous to perform on public roads.

Aurora has access to closed-course test sites that enable us to efficiently test specific software on our vehicles in a broad range of environments and conditions. Tests are allocated to the track that best represents the elements of the ODD most relevant to the objective of the test. Some tracks are commonly used for nominal highway testing as they contain multiple lanes of travel and support high-speed operation. Other test tracks have traffic lights and surface street infrastructure elements that make them well suited for testing intersection interactions, including with vulnerable road users.

Track testing allows Aurora to exercise the Aurora Driver in portions of the ODD that are not commonly encountered during public road testing. For example, vulnerable road users are rare on the highway where the Aurora Driver spends much of its time operating, but test tracks help us ensure we have high-quality data and training for such encounters. Consider the graphic below, which illustrates the pedestrians we have encountered in different parts of the roadway (offroad, near a vehicle on the shoulder, and in-lane).

As pedestrians go from standing to crouching to lying down, not only does it become harder to see them, but the frequency of observing this scenario becomes more rare. Since we know we need to be prepared for these encounters, we need to generate the data to train the Aurora Driver to handle these situations.

Graphic. Vulnerable road users in various positions across the roadway

This is where our test tracks come into play. With a closed test track, we can safely collect data on rare encounters, including pedestrians lying down on the road in various positions (e.g., parallel, perpendicular, or diagonal to the lane) to create large datasets. Providing more data and examples benefits the Aurora Driver and helps improve on-road performance.

1.4.4

On-road testing

The Aurora Driver operates physical trucks in the real world and, while we can simulate what other actors might do in various situations, it is important to observe how public road drivers actually interact with our vehicles. For example, will other road users understand the Aurora Driver's intentions when merging? Driving on public roads autonomously allows us to gather data on such everyday situations.

On-road testing also helps us continually create and refine virtual tests. Our engineers incorporate meaningful events flagged by vehicle operators, as well as those automatically flagged by the Aurora Driver itself, to generate virtual variations of on-road events for virtual testing permutations that can be continually used to improve and fine-tune existing Aurora Driver capabilities. These virtual testing scenarios are based on two on-road testing sources:

Annotations

While hauling freight for pilot customers and conducting testing, vehicle operators monitor the vehicle's performance and flag scenarios that may provide the Aurora Driver with learning opportunities. The detection of many of these scenarios is also automated. We then recreate these scenarios in our Virtual Testing Suite to prepare the Aurora Driver for a diverse set of situations on the road.

Vehicle operator interventions

Our vehicle operators are trained to engage and disengage the self-driving system when appropriate. A vehicle operator may proactively disengage the Aurora Driver and retake manual control of the vehicle when they believe there is a chance that an unsafe situation might occur or they believe the vehicle was not behaving as designed. Each intervention is flagged and may be replayed later in virtual testing to determine how well the Aurora Driver would have handled the situation had our vehicle operator not taken control. Ultimately, these interventions are used to continually improve the Aurora Driver.

1.5

External HMI

As part of being a responsible road actor and in order to communicate with other drivers, our Class 8 trucks are equipped with an external human machine interface (HMI) in the form of illuminatable lights mounted on the exterior of the vehicle.

The HMI indicates the Aurora Driver's state in a comprehensible manner to external actors, and is capable of illuminating various colors in steady-burn or specific flash patterns to indicate different meanings to external actors. For example, the HMI's activation of steady-burn white lights will indicate the Aurora Driver's autonomy state to Aurora team members within our terminal, while flashing amber lights will warn approaching traffic of the truck's presence if stopped on the shoulder of a highway.

Figure. External human machine interface

1.6

Safety of the vehicle platform

The safety of the base vehicle platform is key to any road-worthy self-driving vehicle.

We have partnered with established original equipment manufacturers (OEMs) that have the experience and expertise to deliver vehicles with the required steering, braking, and safety features mandated by U.S. regulations to ensure safe operation. We are working closely with these OEMs to co-develop autonomy-enabled trucking platforms—leveraging both their safety expertise and our own.

Aurora is integrating the Aurora Driver into vehicle platforms that have already been certified to meet or exceed applicable Federal Motor Vehicle Safety Standards (FMVSS). We work closely with our OEM partners when integrating the Aurora Driver, co-developing technical requirements and interfaces between the vehicle platform and the Aurora Driver.

The experience we share with our OEM partners directly informs the development of their future vehicles, much like the experience they share with us informs the development of the Aurora Driver. We will continue to work with our partners on testing, development, and strengthening our understanding of how these customized vehicles and the Aurora Driver can seamlessly work together.

1.7

Release process decisionmaking

1.7.1

Risk evaluation

Aurora's Safety Review Board, a group of senior leaders in the company who are responsible for risk and safety management, regularly convenes and conducts a Global Risk Review, which is the culmination of the risk lifecycle.

During the Global Risk Review, the Safety Review Board examines outstanding risks in detail to understand the contents and to calibrate overall risk tolerance. This process ensures that Aurora's senior leadership is able to maintain awareness and acceptance of the overall safety risk profile.

An important part of Aurora's safety strategy is how the decision is made to *launch into an ODD*. This decision is made through Aurora's internal release process. As Aurora's ODD expands, new routes are added, new software capabilities are added, and enhancements are made to its platforms, Aurora will follow a systematic process to ensure its safety case remains closed, and business and legal risk considerations are considered.

Using the existing SMS structure, Aurora has designated process owners who review the status of its safety case, the results of its verification and validation testing, the status of key metrics, supporting risk assessments, and the readiness of supporting operational procedures. Aurora's Safety Review Board reviews this supporting data, and will collectively make a launch decision based on a holistic assessment of safety. These decisions are governed by Aurora's release process and are based on specialized criteria from functional divisions, including software, hardware, security, operations, and safety.

At the time of any driverless launch, we will have established an acceptable level of performance. Aurora's investment in robust metrics helps us measure our actual performance against our performance bar at the time of a release process decision. If our performance is acceptable, then we are ready for launch. A key threshold is the completion of the safety case before we can commence driverless operations on public roads. Because each release process is tied to a specific ODD, we are able to precisely calibrate our bar for acceptable performance. Finally, Aurora's Safety Review Board reviews every release process decision. If there are outstanding issues, the Safety Review Board may request additional work to remedy the issue before granting its final sign-off.

As discussed in this report, Aurora outlines safety requirements that are tested to ensure proper performance. If the requirement passes, we update the evidence associated with that requirement (or claim), and then the claim is closed. If the requirement does not pass testing, the risk is identified and then discussed and either resolved, mitigated, or ultimately accepted by the Safety Review Board.

02 | Fail Safely

2.1

Responding safely to faults

Regardless of how diligently we work to build durable, robust, reliable systems, we must anticipate that system failures will occur. Aurora's Fault Management System (FMS) allows us to continuously monitor, rapidly detect, and safely respond to off-nominal situations.

Our fault management framework monitors and reports on the health of internal diagnostic checks across the whole system, maintaining a virtual checklist of items to track. The framework is designed to constantly monitor a range of factors that contribute to the health of the autonomous vehicle and detect when an issue arises.

When failures occur, the fault management framework maps individual item failures to requirements-driven, predetermined fault responses, initiating the appropriate fault response and preventing autonomy engagement when appropriate.

We developed the Aurora Driver's FMS such that the vehicle can detect faults and operate safely when something goes wrong. It achieves this by performing a minimal risk maneuver[5] such as pulling over to achieve a minimal risk condition[6] by being outside the flow of traffic.

In the event of a loss of a function or component, the Aurora Driver-powered trucks are capable of reaching a minimal risk condition by executing minimal risk maneuvers, because fault tolerance and the ability to fail safely are built into the software, hardware, and embedded systems that make up the Aurora Driver. By employing appropriate segmentation, system integrity validation, and trusted signaling of the FMS, we are able to protect critical safety functions against threats and risks within the ODD. For a fault like a tire blowout, this would mean detecting the damaged tire, evaluating the health of the rest of the system, and pulling over to the side of the road if, when, and how it is safe to do so.

[5] A minimal risk maneuver is a set of vehicle maneuvers intended to reduce risk of harm in the presence of a fault, failure, or other anomaly by attempting to bring the vehicle to a reasonably low risk end state (e.g., pulled over onto the shoulder out of traffic) before the anomaly results in hazardous behavior.

[6] A minimal risk condition is the end-state reached by a minimal risk maneuver or other fault response.

2.2

Fault detection

These health checks correspond to preconditions for nominal execution, and a failed health check indicates a fault has been detected. When a fault is detected, the FMS will evaluate its severity and determine the impact it will have on the Aurora Driver's ability to drive safely. If it is not safe to continue operations, the FMS will initiate a mitigation strategy.

2.3

Achieving the minimal risk condition

When a fault is detected and diagnosed, the FMS will consider the state of the entire system to decide on a sufficiently safe response.

When appropriate, the Aurora Driver's motion planner will plot the trajectory to achieve the minimal risk condition and then the vehicle will execute that strategy according to the vehicle's environment. Once the minimal risk maneuver has been executed, the vehicle will have achieved the minimal risk condition.

To minimize the likelihood of stopping on the shoulder of a busy freeway, the Aurora Driver will, if possible, drive to and stop at a designated "preferred pullover" location. These locations are just off the highway on frontage roads, where our trucks can pull over in such a way that minimizes collision risk, minimizes nuisance to other road traffic, and facilitates recovery if needed. While executing a pullover to a preferred location, if the failure is resolved prior to reaching the location, the Aurora Driver will continue on its route. If the failure is resolved after having come to a stop, the Aurora Driver may return to its route when cleared by a remote assistance specialist in our Command Center.

2.4

Redundancy

Aurora separates, isolates, and incorporates redundant components to independently provide safety-critical functionality.

By following these practices, we are able to ensure that should any one of the critical subsystems fail, there is another subsystem capable of continuing to provide the overall safety functionality. The separate and redundant backups have a very small risk of failing at the same time for the same reason.

Aurora designs its safety critical systems with appropriate redundancy in order to provide functionality even after a fault or failure. These include redundant steering, braking, power management, motion management, energy storage, communication, and computing.

03 | Continuously Improve

3.1

Safety performance indicators

Safety performance indicators are safety metrics that signal potentially elevated risk, and we track meaningful signals on the performance of the Aurora Driver, related systems, and supporting processes. At Aurora, we use safety performance indicators to understand not only the performance of the Aurora Driver, but also our business operations and how safety is prioritized within our organization itself. We regularly monitor and analyze multiple metrics to better understand our processes and identify potential areas for continuous improvement.

Our approach to safety metrics was informed by key industry best practices (see the Be Trustworthy section for additional information). Ultimately the development and usage of our specific safety performance indicators play a critical role in our safety case as they provide evidence to support numerous claims. At the system level, Aurora's safety metrics and specific safety performance indicators are computed using data collected from all of our testing and development activities.

High-level examples of the Aurora Driver data we may utilize in the formulation of our safety metrics include:

- **Scenario test results**
- **Execution of behavioral competencies**
- **Vehicle dynamics (velocity, jerk, acceleration)**
- **Spatial separation from other road users**
- **System failures, faults, and responses**
- **System reaction times**

In addition, at the organization or operational level we can consider data such as:

- **Property damage**
- **Inspection and maintenance results**
- **Execution effectiveness**
- **Safety culture measures**

Hard braking is a good example of a safety performance indicator that many drivers inherently understand. An instance of hard braking is not necessarily a bad thing—after all, if a driver needs to brake hard due to a surprising traffic condition to avoid a crash, then that's a successful outcome. What is more relevant is if the rate of hard braking begins to increase, or deviates, from historical expectations. This may indicate that something has changed and would prompt an investigation to understand the root cause. Ultimately, the goal of monitoring indicators like these is to identify and resolve underlying issues before they become hazardous events in operation.

3.2

Configuration and release management

3.2.1

Configuration management

The goal of configuration management is to enable reproducibility, traceability, and diagnosability, and Aurora achieves this goal through versioning and extensive process controls for various software and hardware assets.

While all software assets are versioned by the build system, a lot of assets also have human-readable versioning schemes, such as an on-vehicle software release and a vehicle driving log. Configuration management also applies to third-party assets, including commercial off-the-shelf sensor firmware, vehicle platform electronic control unit hardware and software versions, and third-party libraries and tools. Collectively, Aurora's configuration management ensures that only tested and approved changes are incorporated into the software builds and that the correct version of the Aurora Driver is deployed on the right hardware for the corresponding vehicle platform.

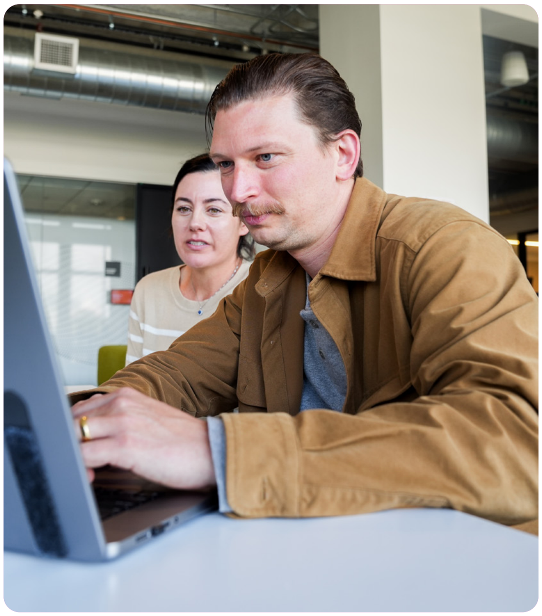

3.2.2

Release management

Aurora prioritizes a well-defined release process that provides high confidence in safety and product capabilities while also meeting velocity and cost expectations. The release process is cross-functional and relies heavily on tools and processes built to accommodate software build, verification/validation, deployment, and analysis. With continuous integration and continuous delivery, verification gates and checks enable requirements to be met for commercial readiness.

04 | Operate with Resilience

4.1

Cybersecurity

Securing an autonomous vehicle requires diligence throughout its development and operation.

A secure system is one that minimizes architectural weaknesses and is ready to respond and recover from identified risks.[7] An ADS's security architecture must also achieve these goals in tandem with operational and safety functions. Aurora's security architectural approaches are motivated and measured through integration into Aurora's safety case. Similar to the safety case, the security development and assessment processes incorporate the ODD as a core constraint in identifying threats.

Leaning on the safety case and security principles, Aurora has developed an extensive and adaptive security approach, aligned with best practices and standards, to secure the varied component ecosystems that compose an autonomous system. We consider all functional areas of our technology to be potential targets with different threat models, and, therefore, potential vehicle safety concerns.

Aurora has adopted security architectures and risk-based assessment methodologies that derive and measure security controls through two major themes —"Trust the Operation of the Aurora Driver" and "Detect, Respond, and Recover". These two major themes are comprised of six approaches that we address cross-functionally with our partners and across the company.

These approaches, and the controls they inform, serve as a blueprint for the components that must be assessed, along with the relative depth for each. Inspired by guidance from the National Institute of Standards and Technology (NIST)[8], NHTSA, and industry groups, this method enables Aurora to address security from both a product and a process perspective, and provides defense in depth through layered controls.

[7] Ross, Ron, Mark Winstead, Michael McEvilley. Engineering Trustworthy Secure Systems. NIST: SP 800-160 Vol. 1 Rev. 1, November, 2022. https://csrc.nist.gov/publications/detail/sp/800-160/vol-1-rev-1/final.

[8] NIST Special Publication 800-160 v2 Developing Cyber-Resilient Systems: A Systems Security Engineering Approach, 2021; ISO 21434 Road vehicle—Cybersecurity engineering, 2021; NHTSA Cybersecurity Best Practices for the Safety of Modern Vehicles, 2022.

4.1.1

Trust the Operation of the Aurora Driver

The Trust the Operation of the Aurora Driver theme is composed of four approaches:

01 Build, Deploy, and Activate
encompasses how we build, test, release, deploy, and activate software for the Aurora Driver platform. This approach drives security controls across the enterprise and onboard ecosystem as it establishes the initial trust of the system. Within the enterprise, Aurora incorporates integrity-protecting security controls such as software signing throughout the software and firmware build and release processes. We employ concepts within our onboard systems such as trusted software and firmware activation where updates are cryptographically validated.

02 Trusted Startup
encompasses diverse security controls supporting the trust of safety critical components and enabling this trust to extend and be evaluated by the autonomy software. This includes booting the computer and the fault tolerant controller's firmware into a trusted state.

03 Engage Autonomy
encompasses the various processes and checks needed to activate and engage the Aurora Driver. We use a trusted software process that is controlled by the Aurora Driver's onboard and offboard systems.

04 Trusted Offboard Actions
refers to secure communication between the Aurora Driver and remote assistance specialists and services. Our security controls ensure that these interactions employ best-practice secure protocols to protect the channel and transmitted data from interception and modification. Our fleet managers, remote assistance specialists, and vehicle operators serve as an important line of defense in mitigating security threats. These professionals are trained to detect, annotate, diagnose, and escalate any potential irregularities in system and vehicle performance.

4.1.2

Detect, Respond, and Recover

The Detect, Respond, and Recover theme focuses on strategic security risk-reducing architectural design and includes two foundational approaches:

01 Identifying Privileged Access
focuses on identifying the necessary authorization needed to support the developmental and operational needs of the Aurora Driver. Credential misuse, over-privilege, and escalation of privilege are key security concerns for any organization, so limiting broad access and enforcing least privileged access may significantly reduce exposure. Furthermore, by identifying and managing privileged access, we can significantly increase the difficulty for opportunistic attackers who would attempt to compromise individual elements and leverage those attacks to further impact the fleet. Finally, vehicles preparing to launch in a driverless mode should only permit privileged and elevated access to a strictly controlled list of identified and approved individuals.

02 Security Detection and Response
focuses on ensuring that a security incident can be detected (both onboard and offboard), and that we can respond to it appropriately and quickly. Detection and response of such incidents are critical functions for an autonomous vehicle, and Aurora has integrated cybersecurity detection engines directly into the fault handling processes of the Aurora Driver. If a detection warrants a response, the appropriate FMS actions are triggered. (Please refer to section 2 Fail Safely on fault management for a more detailed description of how the Aurora Driver achieves a minimal risk condition.)

We're designing our product, our processes, and our tools to continuously and confidently identify the threats, gaps, and risks associated with the Aurora Driver. A significant output of this effort is a comprehensive Threat Assessment and Risk Analysis (TARA) of the Aurora Driver ecosystem. A TARA identifies, justifies, prioritizes, and scopes the necessary security concepts and downstream requirements while also completing appropriately scoped assessments demonstrating the Aurora Driver's resilience to ODD cybersecurity threats. This activity, along with periodic external assessments, drives the mitigation efforts to address issues and reduce risk to an acceptable and appropriate level, while also verifying and validating the overall security posture of the Aurora Driver.

4.2

Incident response

As part of our efforts to improve safety on the road for everyone, the Aurora Driver is designed to minimize the chance of a collision. When the Aurora Driver detects that a collision may occur, it will initially attempt to guard against a potential collision, such as by nudging away from an incoming actor or changing speed. The Aurora Driver may take more evasive actions to avoid a collision or to minimize the impact if a collision is unavoidable. In scenarios where impact with an actor or object occurs, the Aurora Driver is designed to detect the collision and command a post-collision behavior. Additionally, upon impact, the Aurora Driver will notify a remote assistance specialist that a collision was detected. Finally, once the vehicle comes to a complete stop, the Aurora Driver will enter an immobilization state to prevent any further motion that may present a hazard for first responders, bystanders, or involved actors.

The graphic below (to be viewed from left to right, with the left indicating the time before a collision) provides an illustration of the Aurora Driver's actions.

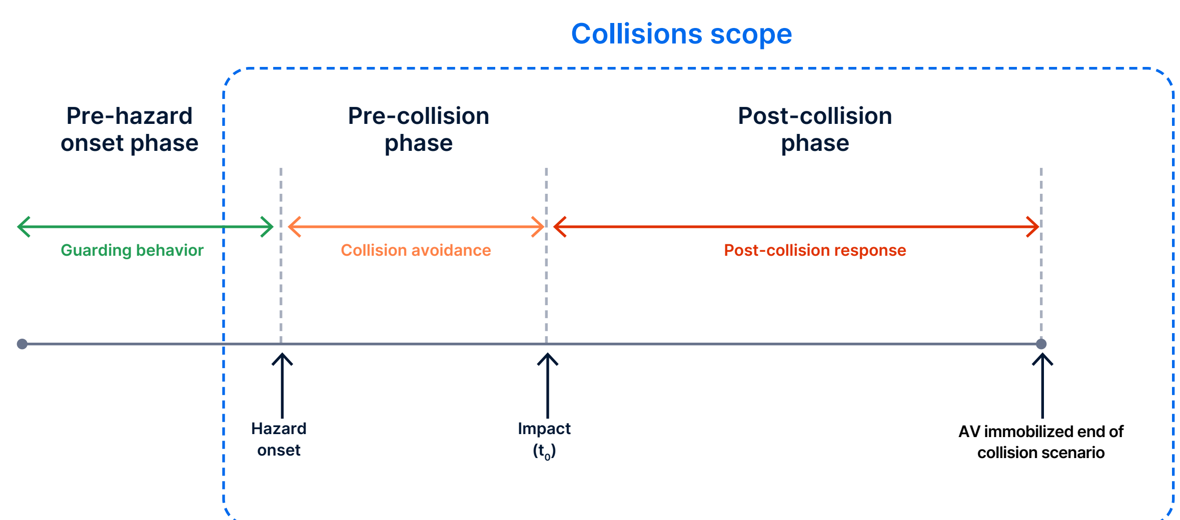

Figure. Collision detection & response

In the event of an incident, Aurora will initiate a series of post-crash procedures, which include communicating with first responders and deploying representatives from our response teams to the location. We are respectful of and reliant on the important incident response role played by public safety officials, including law enforcement, fire and rescue officials, and emergency medical technicians. We proactively engage with public safety officials in the locations where we operate to ensure that our incident response procedures are complementary to those of public safety agencies. This includes instructing first responders on how to determine that the Aurora Driver is immobilized by referring to the external HMI. When the lights in the sensor pod above the side mirrors asynchronously flash an amber light, this means first responders are clear to approach the vehicle so, for example, they can call our 1-800 number and access the document lockbox.

4.2.1

First responder safety interactions

Aurora meets regularly with public safety officials to provide information about how to respond to emergency situations involving the Aurora Driver. As part of this effort, Aurora has published a Public Safety Official and First Responder Interaction Plan, which was developed with input from these stakeholders. It is publicly available here.

The plan provides first responders with details about the truck, including how to: identify the truck, determine if the vehicle is immobilized, access documentation, and contact Aurora support. The contents of the guide are aligned with applicable guidance from the AVSC's Best Practice for First Responder Interactions with Fleet-Managed Automated Driving System-Dedicated Vehicles (ADS-DVs).

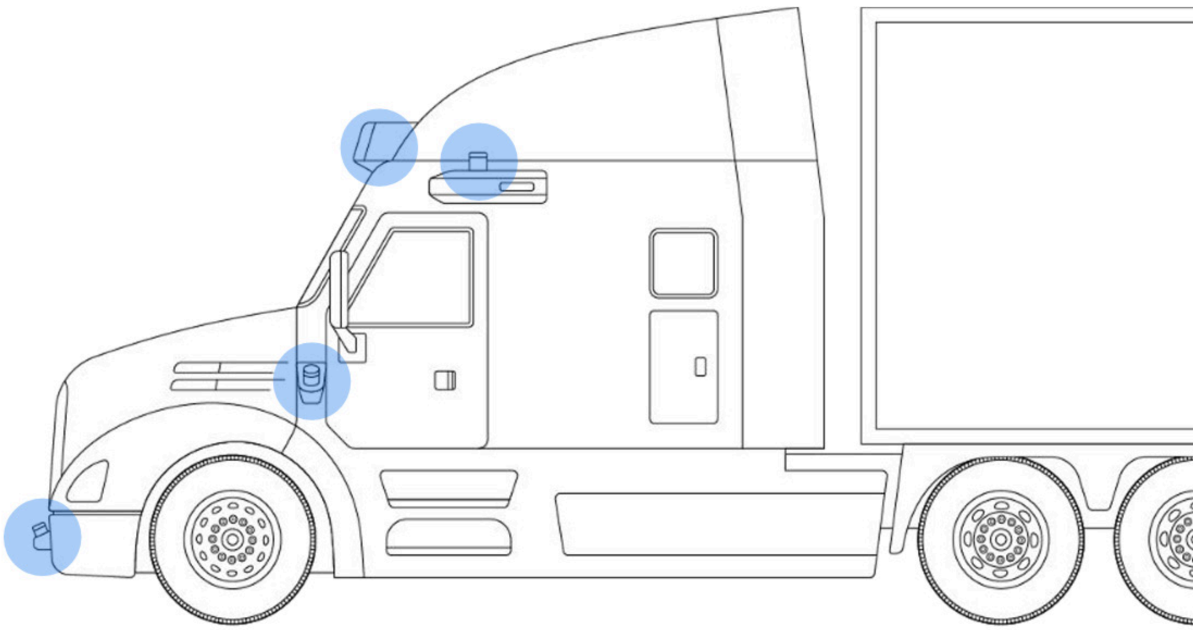

4.2.2

Incident response planning/preparation

Broadly speaking, Aurora's approach to incident response is represented in the workflow diagram below:

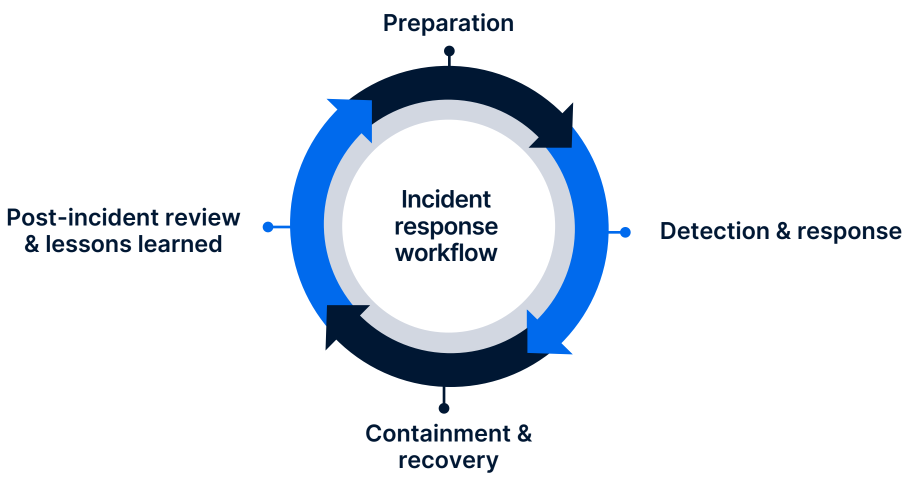

Figure. Incident response workflow

Preparation
We develop a plan that clearly defines roles and responsibilities and decisionmaking authority across the organization.

Detection & response
We outline how we confirm that an incident has occurred, thus activating the incident response plan, including contacting first responders.

Containment & recovery
We contain the incident and take steps towards recovery, including addressing the root cause of the incident and restoring systems to normal operation.

Post-incident review & lessons learned
We seek opportunities to learn and improve. These lessons learned should be fed back into the incident response plan for future use.

Aurora's incident response plan is captured in a comprehensive document that contains the steps to be followed by each relevant Aurora team in the event of an incident. A dedicated Incident Commander is responsible for managing the incident and has authority to direct and approve actions pertaining to the incident. Aurora regularly exercises its incident response through tabletop exercises to ensure that the program is effectively executed and the organization is prepared.

4.2.3

Data recording

The Aurora Driver features a data-logging system which stores vehicle sensor information, the operating state of the self-driving system, and vehicle performance information. We have designed the logging system to secure vehicle log data in the event of a crash. If a crash occurs, the data-logging system stores pre-defined data from the vehicle that allows us to reconstruct the event, as well as provide the data to relevant government agencies, when required.

4.3

Command Center

The Aurora Driver may encounter a variety of situations on the road that benefit from a human's judgment or support, such as speaking with law enforcement officials if necessary or advising on the appropriate course of action in the event of an unusual road closure.

To help improve the Aurora Driver's ability to manage these complex interactions, Aurora designed a software platform to help fleet managers realize maximal utilization of their autonomous vehicles.

Aurora operates a Command Center, which is responsible for all activities related to the management of our fleet of autonomous vehicles. The Command Center performs a number of specific functions, such as dispatching vehicles and asset management.

It is actively staffed at all times when our vehicles are on the road. Aurora's autonomous vehicles are equipped to communicate their health status, location, and observations or sightings on the road with the Command Center, whose specialists may initiate a support session should it be needed.

4.3.1

Remote assistance

Aurora's remote assistance specialists staff the Command Center and remotely monitor and provide information or advice to a fleet of vehicles. They have access to the vehicle's sensor feeds and, similar to the way traditional dispatchers do for human drivers today, they can provide high-level suggestions to an individual vehicle that will enable it to respond appropriately to the situation.

Importantly, this fleet management platform does not allow for remote driving of the truck; the Aurora Driver remains in control of the vehicle and will act on a remote assistance specialist's suggestions only if and when it can do so safely.

4.4

Foreseeable misuse

Aurora defines misuse as usage in a way that was not intended, and organizes foreseeable misuse into two categories: usability errors and system design errors.

A usability error occurs when there is a mismatch between the design intent and a user's expectations, abilities, or needs. Usability errors can be caused by lack of experience, oversight, or fatigue. A system design error occurs when there is a mistake or flaw in the system design or the implementation of a feature.

Aurora employs different mitigations to reduce potential usability error and system design error misuse. For example, to mitigate potential usability error, Aurora has developed clear Standard Operating Procedures, training, and continuous education to guide operations personnel and role-based access controls for critical systems so that only authorized users can access the application needed to launch a trip in autonomy. To mitigate potential system design error, Aurora's continuous verification and validation of system performance allows us to detect system design limitations early and implement enhancements as needed.

05 | Be Trustworthy

5.1

Safety Management System

Aurora's approach to enterprise safety is deeply rooted in our implementation of a Safety Management System (SMS), which is an organizational approach—employed by multiple safety-critical industries, like commercial aviation and nuclear energy—that builds processes to make companywide decisions based on safety risk.

This approach ensures that safety information is presented to the right person, at the right time, and that there is accountability and transparency for every safety action taken across the company.

Aurora built our SMS program on the long-standing guidance from the aerospace industry, tailored to the autonomous vehicle space. This constitutes four "pillars"—Safety Risk Management structure, Safety Assurance program, Safety Policy documentation, and Safety Promotion.

When all these pillars have healthy programs, they help to foster an overall positive Safety Culture. Our SMS helps us proactively identify safety issues and resolve them as quickly as possible. Aurora's SMS conforms with the AVSC Information Report for Adapting a SMS for ADS and underwent an external audit from TÜV SÜD for validation, as described below.

5.1.1

Safety policy

Aurora maintains numerous safety policies which include, but are not limited to:

Aurora's approach to organizational safety

This document establishes expectations for Aurora's SMS, and how all four pillars function across the organization. Aurora's organizational safety approach establishes frameworks for risk-based decisionmaking with delineated roles, responsibilities, accountability, and authority.

Defined risk owners and safety accountabilities

Ultimate accountability for safety at Aurora rests with the Accountable Executive—our CEO, who is supported by the Safety Review Board, comprised of the top leaders of functional areas within the company (e.g., Software, Hardware, Operations) called "process owners". Process owners have ultimate responsibility for risks derived from their functions and accountability to remedy those risks.

Top level safety objectives to maintain and close its safety case

The safety of Aurora's product and operations is governed by the safety objective to create, complete, and maintain a closed safety case for each product release to public roads, including the ODD where that product will operate. Aurora's safety cases are scoped at the enterprise level. This means that, in addition to specific claims related to demonstrating that the product is acceptably safe to operate within its ODD, there are also expectations at the company level to support continuous improvement and maintain transparency with external stakeholders.

The safety case and SMS are integrated, with specific claims related to SMS policy and processes. Aurora uses all elements of its SMS (including the Safety Review Board and assurance programs) to manage its implementation and management of its safety case.

Safety first policy

This policy is maintained within the Aurora company handbook and lays out our non-retaliatory safety reporting policy. This means that the act of submitting a safety concern will not be used for disciplinary or punitive action.

Grounding policy

This policy outlines the procedures for grounding and ungrounding our fleet of test vehicles. The policy is intended to guide the local, organizational, and operational response to grounding events, including incidents or imminent issues involving Aurora vehicles. Additionally, it outlines the identification and resolution of the systemic/engineering issues that may cause such incidents. Critically, Aurora employees are empowered to report issues for potential grounding of the fleet if they believe an unreasonable risk to safety exists.

5.1.2

Safety assurance

Aurora has developed multiple tools to continuously monitor its safety performance, including:

Safety concern reporting

Safety Concern Reports are a crucial element of our SMS and provide valuable information on safety. Reports can come from employees, contractors, or vendors and, once one is received, the Safety Assurance team works with relevant internal subject matter experts to better understand the safety issue, determine a root cause, and assign any mitigating actions.

Safety investigations

The Safety Organization will conduct a variety of investigations if hazards are identified in the organization, both proactively and reactively. All investigations are conducted following the principles of maintaining a Just Culture and following a defined safety investigation process, detailing root cause and corrective action processes.

Driver monitoring systems

The Aurora autonomous vehicle fleet is outfitted with dashcams, equipping operations management with the ability to track vehicle locations, monitor vehicle diagnostic data, maintain Electronic Logging Device/Hours of Service compliance, and review driver behavior and attentiveness. The devices record applicable vehicle activity and automatically identify various triggered events such as vehicle speed variations, harsh braking events, seat belt usage, distracted driving, and others.

Safety performance indicators

At Aurora, safety performance indicators are data-based safety metrics used for monitoring and assessing safety performance. They are explained in detail in a prior section of this report.

Safety surveys

Each year, safety-specific questions are embedded into Aurora's employee survey. The results of these surveys are carefully analyzed and converted, as applicable, into corrective actions to address.

Audits and inspections

The Safety Organization will lead a variety of audits and inspections for the organization. Audits and inspections are conducted by a trained/qualified auditor using a standardized checklist.

5.1.3

Safety risk management

Aurora's Safety Risk Management processes provide a structured approach to identify, analyze, and manage safety risks efficiently and effectively. This includes using best practices, utilized across safety-critical industries, such as:

Defined policies for when risk assessment is needed

Aurora has integrated defined triggers for risk assessment into its verification and validation and safety case processes, enabling our Safety Review Board to have cogent discussions on overall risk based on the company's stage of product development.

Defined risk owners in the organization

Using the Safety Review Board structure, every risk identified has a pre-defined pipeline of ownership and accountability.

A unified risk assessment process

Aurora uses a common risk matrix that integrates the best practices from the many safety disciplines Aurora depends on.

Defined risk analysis procedures

Aurora has created risk analysis procedures for Aurora's unique model and operating considerations. Aurora's teams share a policy and playbook, which defines the process for risk arguments, data applicability, and follow-up procedures.

5.1.4

Safety promotion

Promoting and instilling confidence in the organization's commitment to safety, and each employee's role in safety, is critical to building a positive safety culture. Participation in Aurora's safety programs is everyone's obligation, and a culture that encourages positive contributions to safety encourages more of the same. Aurora's Safety Promotion efforts help to accomplish goals through:

Onboarding training for all Aurora new hires

Every new hire undergoes safety training in their onboarding. This includes an overview of Aurora's safety approach, the policies that pertain to all employees, how Aurora thinks about Safety Culture, and how to file a Safety Concern Report.

Safety refresher training for all employees

Every employee is periodically required to undergo safety refresher training, which includes the latest safety policies. These courses, some with defined test criteria, are designed to ensure retention and overall training effectiveness.

Direct questions asked in internal employee surveys

As detailed in the Safety Assurance section above, internal surveys are issued to employees on a regular cadence and include direct questions related to safety. This ensures that Aurora leadership is constantly getting feedback on the efficacy of the company's safety efforts, to include the overall perception of safety and organizational climate.

Continual safety touchpoints

Safety-related topics are frequent agenda items in company all-hands and regular internal communications events.

A variety of tailored training based on employee roles

Beyond orientation, a variety of roles in Aurora require specialized training based on job function and the specific hazards that may be encountered. Training is managed within a Learning Management System and periodically audited by the Safety Organization.

5.1.5

Third party validation

In June 2024, Aurora commissioned TÜV SÜD, a global leader in testing, inspection, and certification, to perform a third-party audit of our conformity to the AVSC Information Report on SMS. Over the course of two months, the audit team reviewed documentation and evidence to verify the implementation and effectiveness of all elements discussed above, and performed additional interviews with the CEO, Safety Organization leadership, and a variety of randomized employees throughout the organization to evaluate safety expertise, adherence to core safety programs, prioritization of safety policies, and management of a strong safety culture.

The auditors concluded that the Aurora team goes above and beyond industry guidance, and verified the effectiveness of the overall Safety Management System.

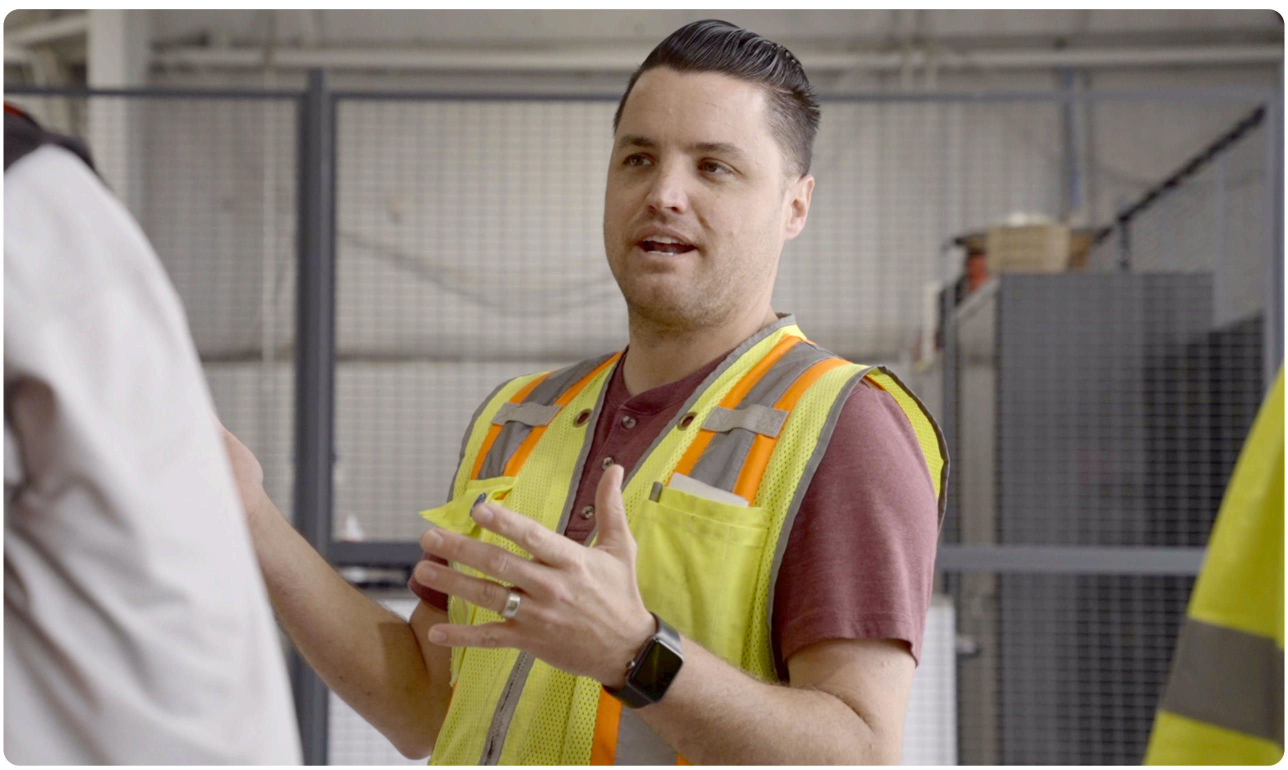

5.2

Safety Advisory Board

Finally, as part of our commitment to Trustworthiness, we frequently engage with stakeholders outside of Aurora, including federal and state regulators and legislators, law enforcement agencies, standards organizations, industry organizations, local public officials, and safety advocacy groups. Through these regular engagements, we keep these external stakeholders apprised of our progress and consider input they provide on our policies and procedures. In 2021, we established an external Safety Advisory Board of esteemed transportation safety leaders, who offer an external perspective on Aurora's overall approach to safety and other topics. Board members have decades of experience in government, academia, and industry, including from regulatory agencies, aviation, trucking, and more.

Aurora meets with this group on a monthly basis and Safety Advisory Board members have experienced rides in the Aurora vehicles on multiple occasions, witnessing our progress firsthand.[9]

Adrian Lund

SAB Chair

Former President of the Insurance Institute for Highway Safety

Karen Rasmussen

SAB Vice Chair

Executive Director of the Independent Carrier Safety Association

Dave Carbaugh

SAB Member

Former Chief Pilot Flight Operations Safety, The Boeing Company

Jeff Runge

SAB Member

Former Administrator of the National Highway Traffic Safety Administration

George Snyder

SAB Member

President and CEO of GHS Aviation Group and previously held key senior executive leadership roles for safety at various commercial airlines

We understand that developing self-driving technology cannot be done in a vacuum and are committed to learning from and contributing to our industry. We believe this highly collaborative approach helps us deliver our products safely and is critical for the responsible development and eventual full-scale operation of self-driving vehicles.

[9] Learn more about our Safety Advisory Board in this blog post https://aurora.tech/blog/aurora-shares-safety-report

5.3

Safety metrics reporting

Safety metrics and safety performance indicators serve additional purposes beyond Aurora's own product development and performance monitoring (see section 3.1 Safety Performance Indicators). Aurora reports certain data, such as crash information, directly to federal and state regulators to satisfy mandated data collection requirements.[10] These data help government stakeholders learn about the real-world performance of self-driving vehicles and inform their regulatory responsibilities.

10 See, e.g., NHTSA Second Amended Standing General Order 2021-01 (April 5, 2023).

5.4

Industry best practices and standards

Aurora believes that being a leader in the development of autonomous vehicles includes collaborating with others in our industry.

That's one reason why we published our Safety Case Framework; to encourage collaboration and promote transparency.[11] We also share our deep safety expertise by participating in working groups and committees across various industry standards development organizations.

The automotive industry has a long history of collaborating on standards development, and we expect collaborative standards development to help shape how self-driving systems are deployed. We've taken an active role in AVSC, SAE, IEEE, and Underwriters Laboratories standards development processes. We anticipate that voluntary industry standards for self-driving technologies will eventually inform the development of specific federal guidance and/or Federal Motor Vehicle Safety Standards.

The U.S. Department of Transportation Automated Vehicles Comprehensive Plan[12] identified a set of twenty automotive best practices and standards which individually address different components or processes for a self-driving system or enterprise. Standards such as SAE J3016, SAE J3018, UL 4600, ISO 26262, and ISO 21448, as well as the best practices from the AVSC, are applicable to self-driving technology.

Safety is an ongoing and iterative process, and as a result, best practices and standards will continue to evolve and mature. We continually incorporate the collective wisdom of the industry in our technology.

[11] Learn more about our Safety Case Framework in this blog post https://aurora.tech/blog/aurora-unveils-first-ever-safety-case-framework

[12] U.S. DOT, 2021, https://www.transportation.gov/sites/dot.gov/files/2021-01/US-DOT_AVCP.pdf.

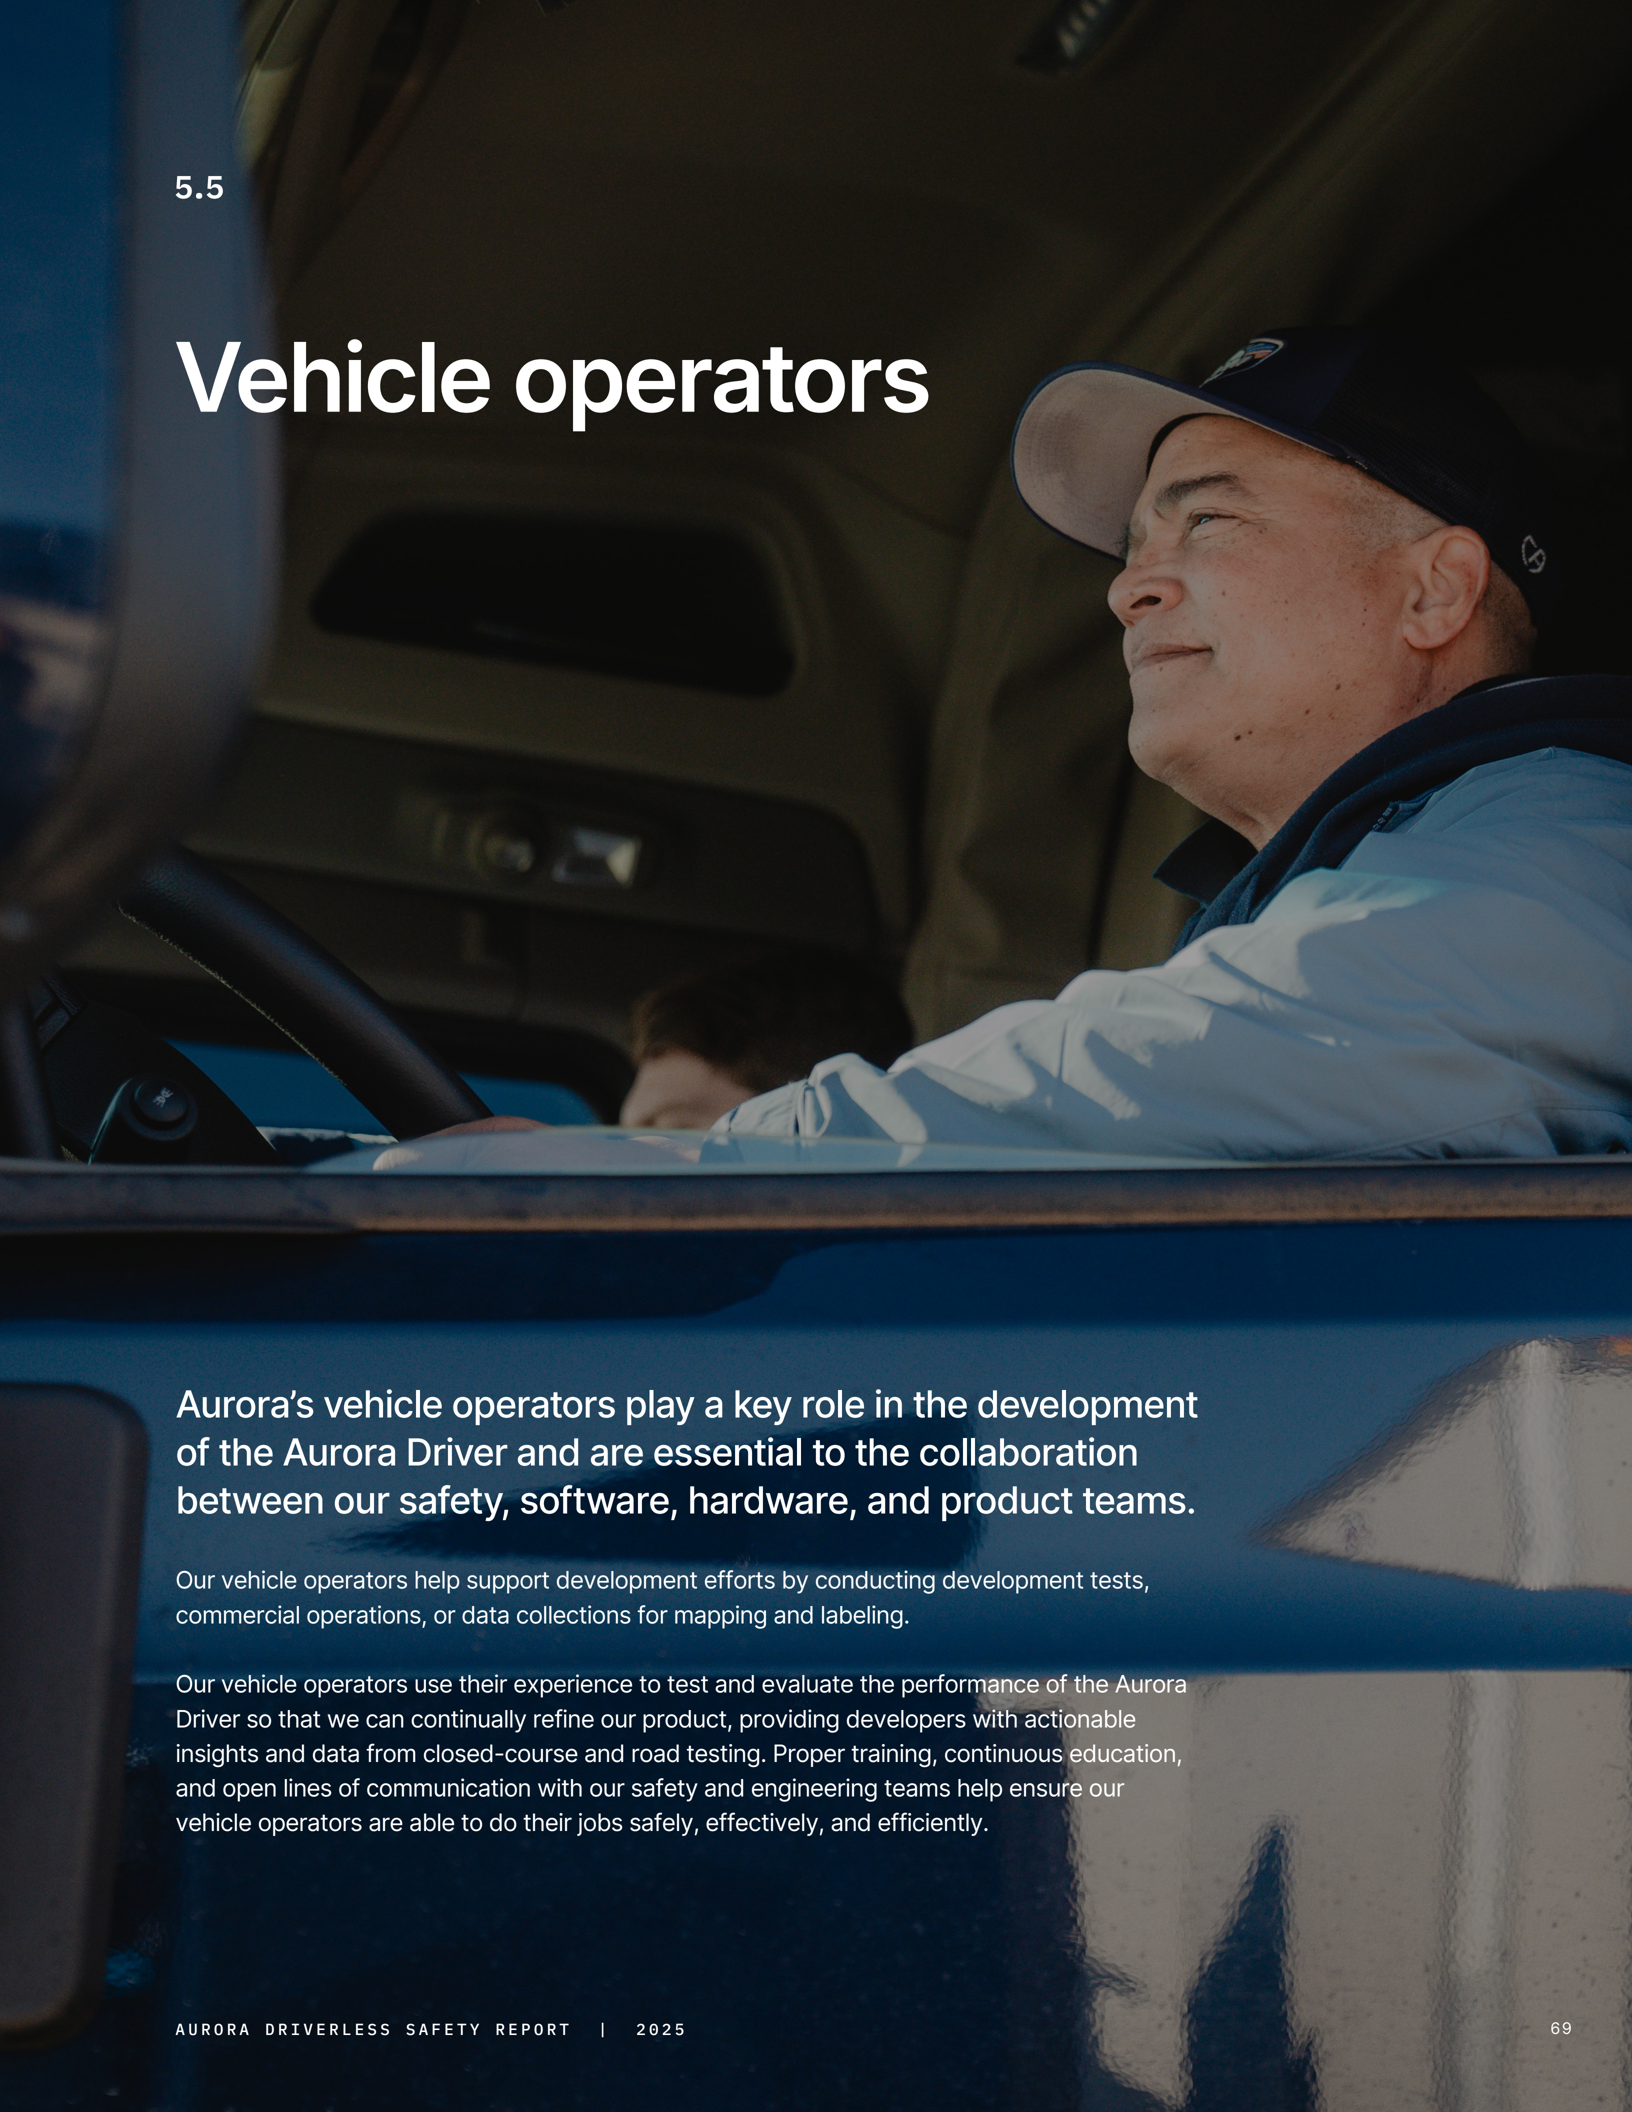

5.5

Vehicle operators

Aurora’s vehicle operators play a key role in the development of the Aurora Driver and are essential to the collaboration between our safety, software, hardware, and product teams.

Our vehicle operators help support development efforts by conducting development tests, commercial operations, or data collections for mapping and labeling.

Our vehicle operators use their experience to test and evaluate the performance of the Aurora Driver so that we can continually refine our product, providing developers with actionable insights and data from closed-course and road testing. Proper training, continuous education, and open lines of communication with our safety and engineering teams help ensure our vehicle operators are able to do their jobs safely, effectively, and efficiently.

5.5.1

Vehicle operator roles in autonomous operations

There are currently two different roles a vehicle operator can play—pilot or co-pilot.

The pilot, holding a valid Commercial Driver's License (CDL), is responsible for keeping the vehicle, the people in the vehicle, and anyone around the vehicle safe at all times. The pilot is able to quickly respond and exert control over the vehicle by actuating the steering wheel or pressing on the throttle or brake pedals to disengage autonomy.

The co-pilot, when present, uses a specialized visualization tool on a laptop to monitor and interact with the Aurora Driver. This tool allows the operator to see a model of the vehicle and how it perceives its surroundings, including other vehicles, nearby pedestrians and cyclists, road lanes, and traffic lights. As autonomous driving is underway, the co-pilot monitors the performance of the Aurora Driver and can inform the pilot of anything that looks out of the ordinary and any potential Aurora Driver performance shortfalls. The co-pilot also annotates testing logs with useful information designed to provide context to notable testing events.

As we prepare for our driverless commercial launch, we are driving some trips autonomously with a single vehicle operator (pilot), meaning there is no co-pilot in the vehicle. We deploy single vehicle operator missions with release candidate versions of the Aurora Driver software when the system's performance is strong enough that a co-pilot is no longer required for oversight or annotation. Co-pilots will continue to support development work as we expand to new lanes, so our fleet will have a mix of driverless, single vehicle operator, and co-piloted vehicles.

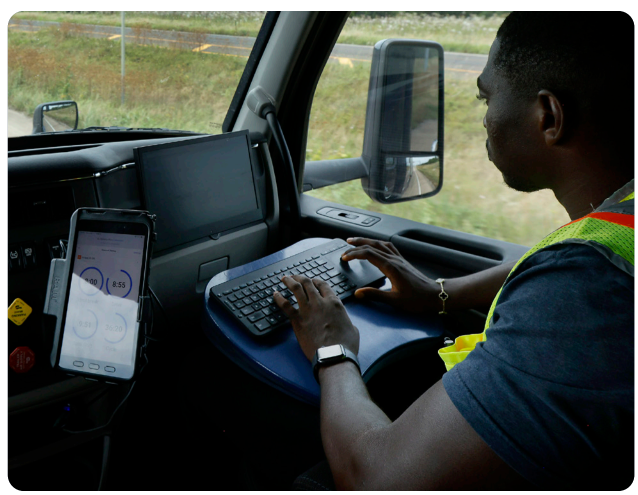

5.5.2

Hiring and training

Our recruiting process for vehicle operators seeks out safe, experienced drivers who have undergone a driving assessment to ensure their ability to operate a motor vehicle in an exemplary manner.

We aim to hire candidates with key attributes like decisiveness, adaptability, awareness, and the ability to think critically and communicate clearly under pressure.

Our vehicle operators are full-time Aurora employees and have diverse backgrounds and experiences (e.g., military veterans, pilots, professional truck drivers, educators), which helps us get varied feedback through the field-testing process. Every one of our vehicle operators has passed a criminal background check, an extensive driving history background check, and a thorough driving evaluation. To ensure that all vehicle operators stay current with new policies, they are required to participate in periodic refresher training programs and mission pre-briefings, which includes material on any new process or procedures.

All pilots possess a valid CDL and have years of experience as a Class A Commercial Driver, and new pilots are enrolled in new hire operator training, as well as continuous education.

Our training program combines hands-on experience inside the vehicle on a closed-course track with classroom learning about core concepts of the Aurora Driver, including operational limitations, ODD operating parameters, vehicle controls, and safety components.

5.5.3

Behavioral policies and technology

We recognize that there are key distinctions between conventional driving and operating a developmental self-driving vehicle, so we have implemented a number of technologies and policies for vehicle operators to maximize the safety of self-driving vehicle operations, which include but are not limited to:

Hours of Service regulations

Our truck pilots adhere to the FMCSA's Hours of Service regulations. Our trucks are equipped with Electronic Logging Devices to ensure compliance with these rules. Additionally, vehicle operators are encouraged to use their own discretion to further limit time operating a vehicle, particularly in conditions that have a higher cognitive demand, such as roads with high pedestrian and bicycle traffic or nighttime driving.

Cell phone and smartwatch use

Pilots are strictly prohibited from interacting with mobile devices and smartwatches while the vehicle is in motion or the self-driving system is engaged.

Monitoring

All of our self-driving vehicles are equipped with a third-party driver monitoring system to help ensure our drivers are not distracted or fatigued.

We have also established a Fatigue Risk Management Program (FRMP) which adopts a data-driven approach to continuously monitor and manage fatigue-related risks within the workplace. Rather than a prescriptive approach, an FRMP adopts a performance-based, iterative approach that creates multiple barriers to protect against fatigue-related incidents. It applies principles from our SMS to achieve a practical balance between safety, productivity, and cost by relying on effective safety reporting culture. Through this program, we can better monitor and mitigate fatigue risk.

While previous prescriptive efforts, such as limits on overtime shifts and conservative weekly hours of service limits, can prevent time-on-task fatigue, they are one-size-fits-all approaches to fatigue management. These efforts do not take into account individual differences between vehicle operators, with factors such as sleep quantity or quality, or task differences that arise within different route types. Therefore, a detailed FRMP aids in the development of risk controls to manage individual and task-related fatigue.

Conclusion

As we have emphasized throughout this self-assessment, safety is the foundation of everything we are building at Aurora.

Developing the Aurora Driver is a complex undertaking. It requires technical expertise, engineering discipline, and a commitment to safety that is imbued throughout all aspects of the endeavor. We believe that we have captured those concepts, and others, in the five principles that guide our work—Ensure Proficiency, Fail Safely, Continuously Improve, Operate with Resilience, and Be Trustworthy. It is through these principles that we are focusing our efforts to deliver the benefits of self-driving technology safely, quickly, and broadly.

From the technology we include on our vehicles to the policies we use to manage our people, we have built a safe product from the ground up. Whether we're designing, engineering, testing, or operating our products and services, we are confident that we have methodically considered and documented our approach in a way that demonstrates both our confidence in and commitment to safety first and foremost.

This document is intended to provide a high-level guide to how we approach safety at Aurora, but it is not comprehensive in the sense that safety is a culture and a mindset—it is constantly under assessment and improvement. This report only provides a brief look into Aurora's engineering and processes, and is just one way that we communicate how we think about safety. We encourage you to stay informed about our latest operations and safety cases by visiting: www.aurora.tech/safety.

Appendix

Aurora VSSA Contents	NHTSA VSSA elements
A letter from the CEO and co-founder	
Introduction	
Executing our safety case	
Aurora Driver for Freight	
01 Ensure Proficiency	
1.1 The Aurora Driver	System safety, OEDR, federal, state, and local laws
1.2 Aurora's Operational Design Domain	ODD
1.3 Safety engineering	System safety, validation methods
1.4 Testing the Aurora Driver	System safety, validation methods
1.5 External HMI	HMI
1.6 Safety of the vehicle platform	Crashworthiness, federal, state, and local laws
1.7 Release process decisionmaking	System safety, validation methods
02 Fail Safely	
2.1 Responding safely to faults	Fallback, post-crash ADS behavior
2.2 Fault detection	Fallback
2.3 Achieving the minimal risk condition	Fallback
2.4 Redundancy	System Safety
03 Continuously Improve	
3.1 Safety performance indicators	Validation methods
3.2 Configuration and release management	Validation methods
04 Operate with Resilience	
4.1 Cybersecurity	Cybersecurity
4.2 Incident response	Data Recording, Post Crash ADS Behavior
4.3 Command Center	HMI, OEDR
4.4 Foreseeable misuse	System safety, validation methods
05 Be Trustworthy	
5.1 Safety Management System	System safety, validation methods
5.2 Safety Advisory Board	System safety
5.3 Safety metrics reporting	Federal, state, and local laws
5.4 Industry best practices and standards	System safety, validation methods
5.5 Vehicle operators	HMI, fallback
Conclusion	

Glossary

ADS	Automated Driving System
AVSC	Automated Vehicle Safety Consortium
CVSA	Commercial Vehicle Safety Alliance
NHTSA	National Highway Traffic Safety Administration
FMCSA	Federal Motor Carrier Safety Administration
FMCW	Frequency Modulated Continuous Wave
FMS	Fault Management System
FMVSS	Federal Motor Vehicle Safety Standards
FRMP	Fatigue Risk Management Program
HD	High Definition
IEEE	Institute of Electrical and Electronics Engineers
ISO	International Organization for Standardization
MRC	Minimal Risk Condition
MRM	Minimal Risk Maneuver
ODD	Operational Design Domain
OEM	Original Equipment Manufacturer
SAE	Society of Automotive Engineers
SMS	Safety Management System
UL	Underwriters Laboratories
VSSA	Voluntary Safety Self-Assessment