Aurora

Safety Report

2022

Table Of Contents

Our mission is to

deliver the benefits of self-driving, safely, quickly, and broadly.

A Letter from CEO and Co-Founder Chris Urmson

At Aurora, we are focused on a single mission: to deliver the benefits of self-driving technology safely, quickly and broadly. Accomplishing this mission requires excellence in many domains: robotics, product development, logistics, public policy and more. **Across all of these areas, we put safety first.**

We're building self-driving products to change for the better the way people and goods move through the world. We work to build a transportation ecosystem that is reliable, efficient, and less expensive. We reject the status quo of over 40,000 annual roadway deaths and refuse to choose between saving lives and delivering value.

We've chosen to take a safety approach that is built around transparency and accountability, and we think about safety throughout our business. As you'll see in this assessment, our five pillars — Proficiency, Fail Safety, Continuous Improvement, Resilience and Trustworthiness — span not just our product, but our development process and our business, and involve engagement with important stakeholders.

Since 2016, the U.S. Department of Transportation has asked that autonomous vehicle companies explain how they prioritize safety through Voluntary Safety Self-Assessments. This transparency is invaluable for earning public trust and, since our early days, Aurora has routinely shared our safety approach via regular public blog posts and updates to our safety self-assessment. This year our self-assessment discusses how we are moving closer to driverless operations and launching our commercial self-driving products through innovations like our Virtual Testing Suite, Fault Management System, and more.

With this report, we're reiterating a crucial pledge to continue to intertwine our safety, development, and deployment workstreams: **Aurora will not launch our autonomous trucking product until our Safety Case for our initial driverless operations is complete.** We see this as the highest safety bar in our industry, and one that helps ensure our complete product (including software, hardware, and data services) and our company, are ready for commercial operations.

Building autonomous vehicles safely is neither simple nor easy. By implementing a thorough, thoughtful safety approach, we set a high standard for ourselves as we define the next century of transportation innovation.

We embrace the responsibility to build this technology safely.

Chris Urmson
CEO & Co-Founder


T680
Aurora

0 | Introduction

Introduction

This report describes how Aurora weaves safety into every part of our development process, from design and virtual testing to operator training and testing on public roads — meeting and exceeding USDOT guidance requirements.

In four consecutive guidance documents, the National Highway Traffic Safety Administration (NHTSA) recognized that, given how rapidly autonomous driving technologies evolve, "no single standard exists by which an entity's methods of considering a safety design element can be measured."[1] As such, NHTSA encouraged entities engaged in automated driving system (ADS) testing and deployment to publish a Voluntary Safety Self-Assessment (VSSA) to demonstrate how they consider and address 12 priority safety design elements recommended by NHTSA.

With this third edition of Aurora's VSSA, we continue to demonstrate how those 12 safety elements are reflected in our overall safety approach. We have also included other elements that we believe provide additional detail on our approach to building and operating safe autonomous vehicles.[2]

We organize our Safety Case Framework, and have organized this report, by five major principles. This 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.

1. NHTSA. Automated Driving Systems 2.0, A Vision for Safety
2. References to Aurora within this document mean Aurora Operations, Inc. and its affiliated entities.

1. Proficient

This principle focuses on our autonomous vehicle's 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 various types of testing (virtual, on track, and on road) to improve the Aurora Driver.

2. Fail Safe

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.

3. Continuously Improving

This principle and section outlines 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.

4. Resilient

This principle focuses on how we've 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.

5. 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.

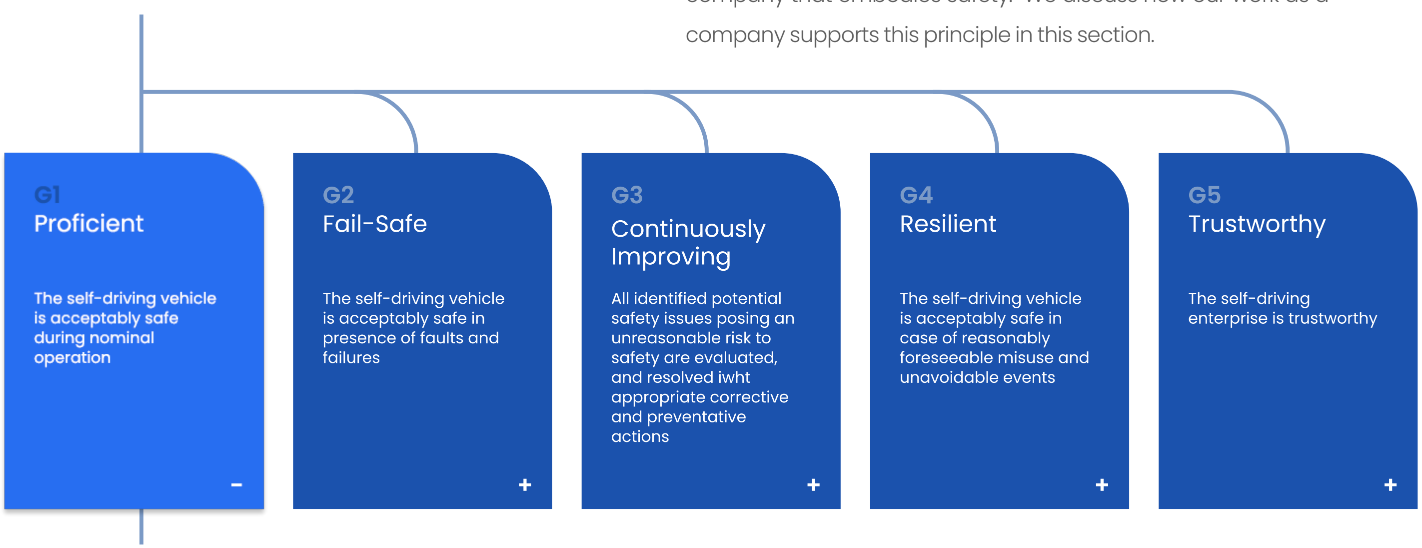

Executing Our Safety Case

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.

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 subclaims, which are then ultimately supported by evidence. We believe a safety case is the most effective and efficient path to safe driverless operations, and an imperative component for any company looking to safely deliver commercial-ready self-driving vehicles at scale.

Safety case 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).[3]

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

For example, as we continue to develop with the Aurora Driver, we currently have vehicle operators in our Class 8 trucks and passenger vehicles monitoring

3. https://aurora.tech/blog/aurora-unveils-first-ever-safety-case-framework. See the Proficient section for an explanation of ODD.

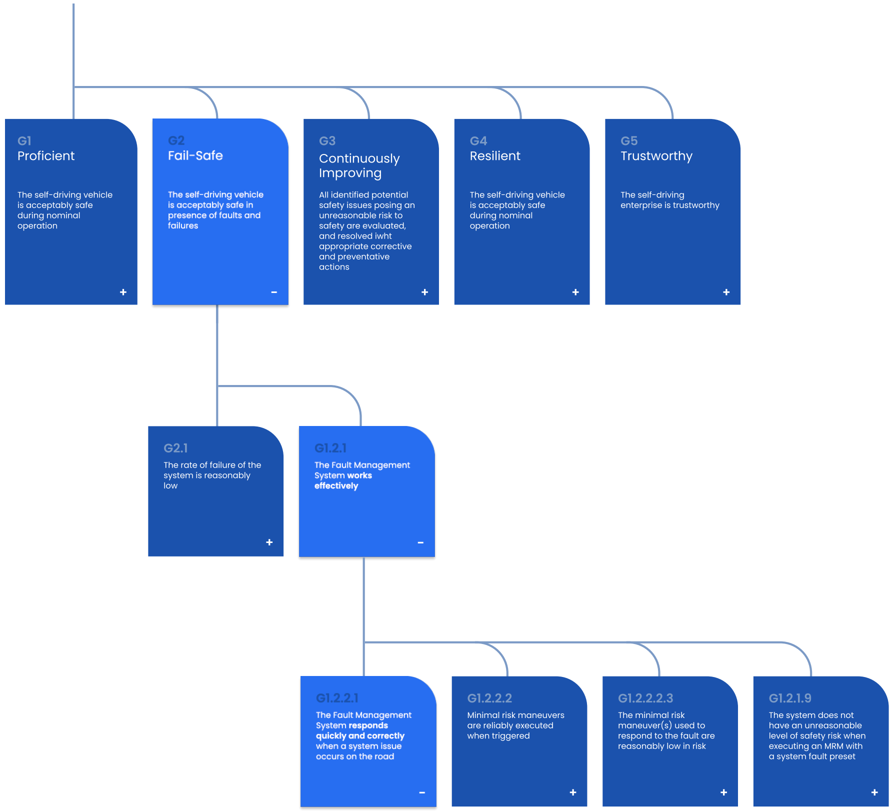

the performance of the Aurora Driver at all times and ready to take over as necessary to ensure operational safety. Therefore our tailored safety case for this use case includes claims focused on vehicle controllability and vehicle operator hiring, training, and operational procedures, among others. However, when we reach the point of removing the vehicle operators from cabs, 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 a completed safety case that requires vehicle operators to be in the cabin to supervise the Aurora Driver. We are actively working to complete the safety case for our initial driverless operations. We will continue to maintain multiple safety cases, each tailored for different use cases as appropriate.

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. Delivering a safe product is a critical part of developing self-driving technology; so is being transparent with our approach. The safety case is one way in which we can show our work.


AURY-5427
Aurora
Aurora
NEXT

1 Proficient

G1
Proficient

The self-driving vehicle is acceptably safe during nominal operation

Aurora's Operational Design Domain

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.

The operational design domain (ODD) refers to the conditions in which the Aurora Driver is intended 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 the Aurora Driver. We then define the ODD based on the conditions, use cases, restrictions, and scenarios within that domain. From there, we map the area and begin simulations of scenarios we expect to encounter within this ODD. Our trained vehicle operators drive the new roads manually to collect data and gain first-hand understanding of the ODD. Our engineers then review the collected data and make any necessary software modifications before the vehicle operators operate the system in self-driving mode within the ODD.

How we operate within the ODD is constrained at both software and operational levels. At the software level, we use geofencing techniques that impose a system prohibition on operating autonomously in areas outside of the ODD. Software constraints can also restrict routing of the Aurora Driver within the geofence based on a set of configurable ODD 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 geofence barrier may not have been created. These may include the review by a remote specialist in response to a vehicle trigger, which will often correspond with a vehicle action to achieve a minimal risk condition.

Operational constraints are developed at the very beginning of our process, starting with product requirements. This ensures that the entire team is aligned on the ODD and what they will need to develop. Before starting a mission, the dispatch team evaluates weather conditions along the route to anticipate whether these

G1.2.1

The Self-Driving Vehicle is **designed for safe operations** within the area it is designed to operate in

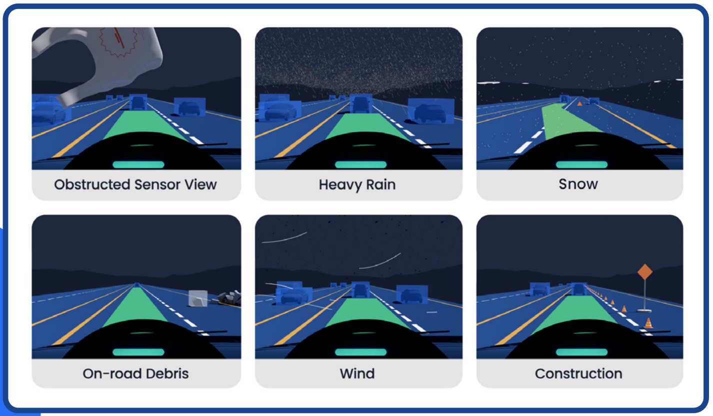

conditions fall outside of the ODD weather parameters. In some cases, we also use our vehicle operators as operational controls. For example, our vehicle operators are trained on the governing ODD characteristics so that they can monitor situations and take manual control of the vehicle when presented with a known condition outside of the Aurora Driver's capabilities.

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

Although we are consistently updating our ODD, as of December 2022, our vehicles operate within the following ODD:

- On public roads in Pennsylvania, California, and Texas;
- 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; and
- In weather conditions for which we understand how the Aurora Driver will perform, including light to medium snow and rain.

We halt testing in autonomy in some of the following circumstances, should we have any concern that the safety of our performance would be degraded. For example:

- Heavy fog, heavy rain, heavy snow, and ice; and
- High wind;
- Adverse road conditions including unplowed roads or those that may result in loss of traction.

The Aurora Driver

We are designing the Aurora Driver so that it can operate with diverse vehicle platforms, using a common core of hardware, software, and data services.

The Aurora Driver consists of the hardware, software, and data services required to safely deploy a self-driving vehicle. We are designing the Aurora Driver so that it can operate with diverse vehicle platforms, using a common core of hardware, software, and data services. This unified technology architecture not only allows us to leverage our development work across multiple use cases — it also simplifies our safety approach by enabling us to focus more on the Aurora Driver.

The Aurora Driver's Main Autonomy Components

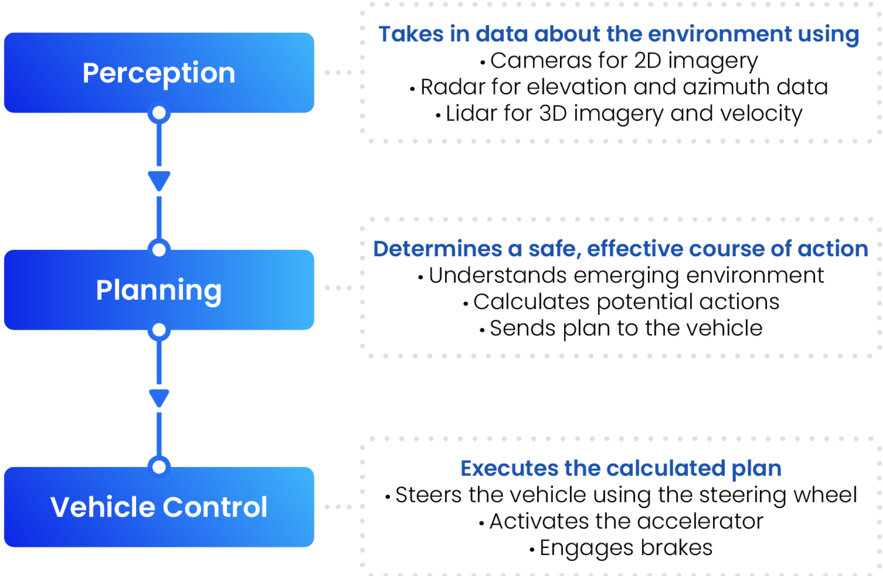

Perception

The first step in driving safely and proficiently is 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 both early and late 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.

Cameras

Cameras observe the visible light reflected off surfaces. We equip the Aurora Driver with multiple cameras with a combination of near-range, medium-range, and long-range cameras distributed around the vehicle. Our cameras' high resolution sensors and quality optics ensure we have the broad range needed 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 have 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.

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, our radars provide overlapping sensor inputs to the Aurora Driver.

Lidar

Lidar systems are a staple of self-driving technology because they create detailed point clouds of the surrounding environment. 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.

Self-driving systems have traditionally been outfitted with Amplitude Modulation (AM) lidars, which work by emitting brief light pulses at a fixed frequency. Lidars can determine the location of objects 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 Frequency Modulated Continuous Wave (FMCW) lidar, named FirstLight, is able to detect and track objects 450 meters away. Compared to traditional AM lidars, which are limited to around 200 meters, FirstLight lidar allows the system to perceive an object nearly 8.5 seconds earlier for vehicles moving at highway speeds. This gives the Aurora Driver a huge advantage — because it can sense farther, it gains more time to react to unexpected obstacles. The FMCW also experiences less interference from the sun or other lidars, because each sensor is designed to respond to only the signals they create, making the data received more relevant and less noisy.

FirstLight's technology also allows us to measure the 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 segment out different actors that are moving at different speeds. Faster object recognition and tracking by the perception system allows the motion planning system to better react and gives 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 more accuracy — a major advantage when it comes to driving at highway speeds.

Planning

Driving is full of complex interactions. The way one behaves on the road affects the way vehicles and other road actors in the environment behave. The Aurora Driver is designed to act safely around others on the road — to be predictable and human-like. We fuse machine learning with formal rules to create a robust planner that smoothly navigates situations while retaining the ability to operate in a safe and predictable manner.

Aurora's proprietary FirstLight lidar, which is able to detect and track objects at longer distances than traditional lidar.

Consider a scenario where a bicyclist hidden behind an oncoming vehicle is suddenly revealed when the oncoming vehicle makes a turn. The Aurora Driver must quickly perceive the cyclist, incorporate it into its forecast, and create a new plan to execute its next maneuver safely. Our approach to prediction and decision-making allows the Aurora Driver to take in this new information and respond appropriately. With the benefits of FirstLight allowing for earlier perception and calculation, the Aurora Driver can plan earlier and more effectively.

Instead of the traditional cascaded approach, where the AV plans its movement after forecasting the movement of other road actors, we take an interleaved approach to forecasting, in which the Aurora Driver predicts how other actors will respond concurrently with planning what it will do. This allows the Aurora Driver to make faster choices as complex scenarios emerge, planning and optimizing a route that will navigate the vehicle in a predictable manner.

Vehicle Control

Once the motion plan has been generated, the Aurora Driver executes the planned actions through the vehicle's steering, accelerator, and brake system by monitoring the current vehicle position and orientation and adjusting the vehicle controls 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 coming in through perception.

Virtually Testing the Aurora Driver

We take a rigorous approach to testing the Aurora Driver. We have heavily invested in our Virtual Testing Suite, allowing us to quickly turn on-road events into multiple opportunities for us to learn and advance. Our virtual testing, combined with closed track and on road testing, enables us to test the Aurora Driver in a comprehensive manner.

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.

We test against scenarios that are specifically designed to help us find edge cases and catch potential errors in our software early, well before that version is uploaded onto our vehicles or used on public roads. When the Aurora Driver has effectively practiced maneuvers millions of times in simulation, we build confidence that it will perform them well on the road. For example, the Aurora Driver performed over 2 million unprotected left turns in simulation before ever attempting one in the real world. Virtual testing is the pillar of our development process because it allows us to:

- **Test more thoroughly.** It's not possible to experience every permutation of potential on-road scenarios in real life. Nor is it feasible to run thousands of tests on the road to evaluate a single change to the codebase. Virtual tests allow us to test many permutations of the same scene in a controlled way.

- **Find and address issues early.** Virtual tests allow us to determine how well the parts of the system are working individually before testing them as part of the integrated suite of tests. This makes it much easier to identify the root cause(s) of potential problems throughout development.

- **Develop and improve capabilities quickly.** Fast, reliable feedback from virtual tests allows us to iterate and improve rapidly without introducing unnecessary safety risks to the public roadways.

- **Develop objective measures of progress.** Each time we make a change to the codebase, we run the same tests with the same parameters. When a new version passes tests that its predecessors failed, we can say it is performing better.

- **Make on-road testing more efficient.** When our software passes thousands of off-road tests, we can be more confident that the Aurora Driver will perform appropriately as those capabilities are exercised on public roads. Virtual testing allows us to streamline our real-world tests to focus on the areas where we need additional real-world feedback.

- **Make on-road testing more impactful.** On-road events can be incredibly valuable learning opportunities. But our approach to on-road testing is not merely about driving more miles — it emphasizes collecting quality real-world data and scenarios and leveraging them within the Virtual Testing Suite. Our triage team reviews interesting events flagged by vehicle operators, then works with engineering teams to identify which ones offer opportunities to improve the Aurora Driver. One real-life event can inspire tens or even hundreds of virtual testing permutations, all of which can be continually used to fine-tune existing Aurora Driver capabilities.

Aurora's Virtual Testing Suite enables extensive, thorough testing of the Aurora Driver by creating many permutations of on-road situations.

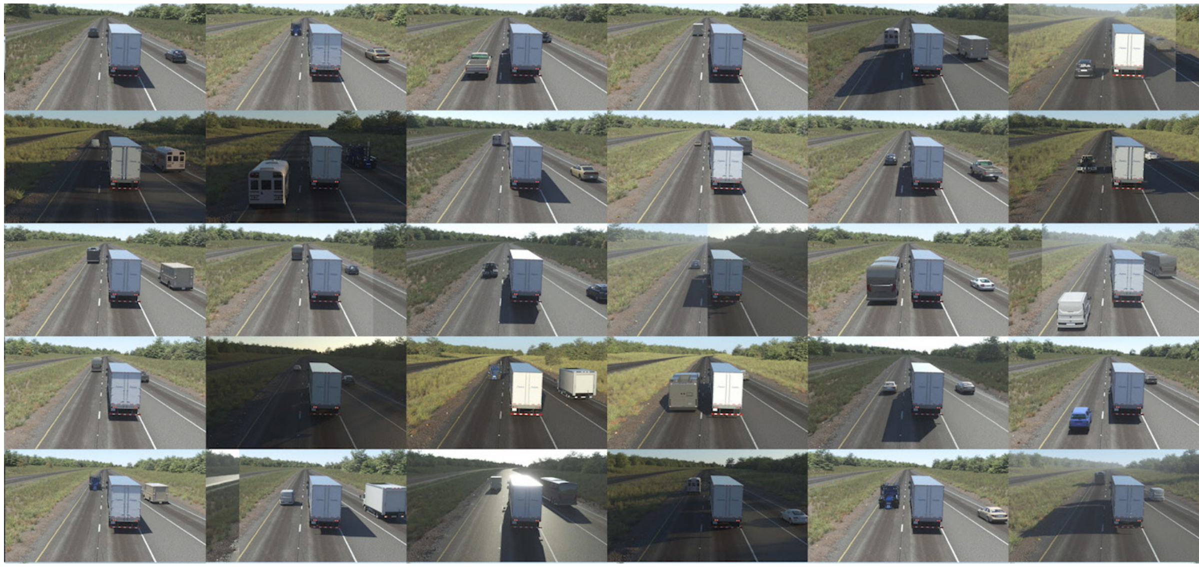

Types of Virtual Tests

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

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 some road debris in the lane of travel. Specialists review footage and then label things like object category (road debris) and velocity (stationary or moving, if it was blowing around). 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.

Manual Driving Evaluations

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 accurately forecast what a trained, expert driver would do in complex situations.

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 the manual miles can provide in improving how we drive autonomously.

Simulations

Simulations are virtual models of the world where we can test how the Aurora Driver reacts in many permutations of the same situation. Simulations also allow us to simulate a wide variety of interactions between the Aurora Driver and other actors in the virtual world.

Examples from our Simulation Database

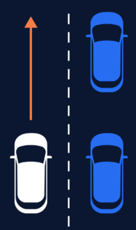

Highway lane keeping

Lane changes

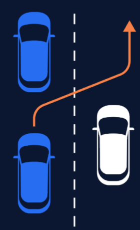

Highway interactions with other vehicles, such as cut-ins and contested lane changes

Correct position in lane

Appropriate speed and distance from other actors

Urban interactions with other actors, such as jaywalkers, cyclists, lane-hogging vehicles

Four-way stop traversal

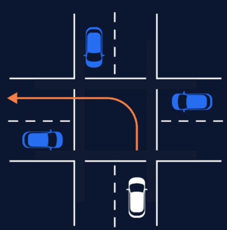

Signalized intersection traversal

Stopping behavior at traffic lights

Go/no-go decision making at amber lights

Passenger comfort metrics

Timely execution of planned maneuvers

Track Testing

We drive the vast majority of our miles in our Virtual Testing Suite, enabling us to make rapid progress. 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.

Aurora has access to a half dozen closed-course test sites to efficiently test specific software on our vehicles in a broad range of environments and conditions. Closed course track testing provides a controlled environment that enables us to observe and understand how the Aurora Driver will perform in the physical world, outside of a virtual test environment. This is especially important because track testing lets us construct and test against situations that are rare or potentially dangerous to perform on public roads.

For example, we could create a scenario in which a vehicle in an adjacent lane drifts into the Aurora Driver's lane, creating a potential collision if the Aurora Driver fails to respond. We cannot control when we may encounter this scenario in real life, but the Aurora Driver needs to be able to recognize and respond to it if it does occur. A closed course test track enables us to safely create, test, and iterate these types of situations to ensure that the Aurora Driver detects this type of encroachment and executes a lane change in response.

Almono Test Track,
Pittsburgh, PA

On-road Testing

On-road tests allow us to assess whether successes in virtual testing translate to the road. We can simulate what other actors might do in various situations, but 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 in autonomy allows us to gather data on such everyday situations with actual users of the road.

On-road testing also helps us continually create new and refine virtual tests. Our engineers incorporate interesting events flagged by vehicle operators to generate additional virtual variations of that event for virtual testing permutations, all of which can be continually used to improve and fine-tune existing Aurora Driver capabilities. These virtual testing scenarios are based on to two on-road testing sources:

- **Annotations.** On each testing mission, one vehicle operator pilots the vehicle, while the other vehicle operator provides support from the passenger's seat, observing the details of the mission, and flagging any scenarios encountered on-road that may provide the Aurora Driver with a unique learning opportunity. 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. 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 opted to take control.

Procedural Scenario Generation

In order to truly reap the benefits of virtual testing and ensure the safety of our vehicles for each operational area, we run thousands or even hundreds of thousands of on-road scenarios, such as another vehicle entering our lane, or failing to stop at a stop sign. But to drive or manually create many permutations of each scenario is impossible, so instead, we have developed systems to automatically generate the massive number of scenarios we need. This process is called procedural scenario generation.

For example, after we specify a location within the Aurora Atlas,[4] the system can determine relevant rules related to maneuver types (e.g., sudden entry into our lane), actor types (e.g., vehicles, bicyclists, or pedestrians), and a range of actor densities (e.g., sparse to dense traffic), and create scenarios that test these permutations.

Procedural scenario generation does not rely on any specific map, so we can change the map via updates or new areas, and then quickly generate new scenarios. Nor does this system rely on any specific vehicle platform, allowing the generated scenarios to be used with any of our existing platform types, including Class 8 trucks and minivans.

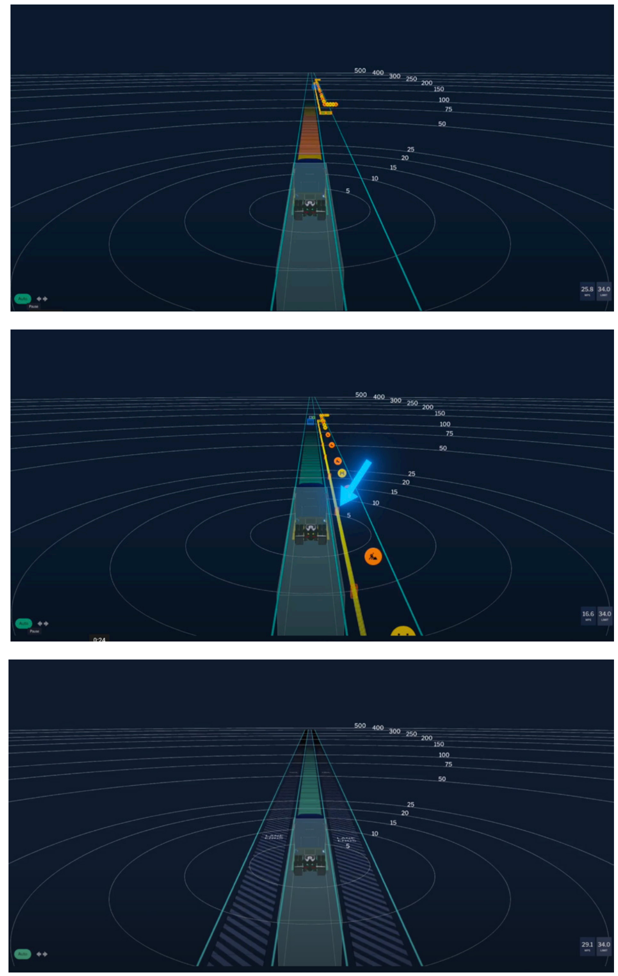

Aurora's Virtual Testing Suite can automatically generate diverse simulation scenarios by changing variables like traffic, on-road actors, and more.

4. Aurora Atlas is our proprietary high definition map for the Aurora Driver and is discussed in greater detail in the "HD mapping and localization" section below.

Vehicle Control Limits

We test the Aurora Driver on public roads in the development phase with trained vehicle operators who are tasked with monitoring the driving situation and disengaging the autonomy system to take over control of the vehicle if they believe the vehicle is performing unexpectedly or potentially creating an unsafe condition. A vehicle operator is able to immediately assume control of the vehicle through the brake, accelerator pedal, or the steering wheel. We also provide a fail-safe disengagement button that will sever control between the vehicle platform and the Aurora Driver.

In addition, we enforce limits on the system's control authority, so that it cannot take actions which vehicle operators may not be able to control. Our vehicle operators are well-trained and extremely capable — but they are still humans, with human reaction times. So these limits lower risk by reducing the allowable range of lateral and longitudinal control inputs sent from the Aurora Driver to the vehicle platform.

Our steering actuation limits are designed to limit vehicle roll-over due to high lateral acceleration, as well as unintended lane departures. Our braking actuation limits are designed to prevent wheel lock up resulting in skidding and potential jackknifing as well as being rear ended as a result of severe and abrupt braking if another vehicle is following too closely. In this way, we can avoid incorrect commands causing a hazard before a vehicle operator could react and correct.

HD 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. Aurora Atlas is our proprietary high definition map for the Aurora Driver. The Aurora Atlas is much more detailed and contains more information than the maps used in typical in-car navigation systems.

By storing precise information about static scenery and road infrastructure in a virtual map, the Aurora Driver does not need to expend as much effort decoding the environment using real-time sensor data. For example, we build stop signs into the Aurora Atlas so the Aurora Driver always knows when to expect them, even when they're faded or hidden behind overgrown vegetation or tree branches. Together, the static scenery and road infrastructure data layers provide data to the Aurora Driver's software system.

We designed the Aurora Atlas not only to 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. A key element of this strategy is keeping the Aurora Atlas locally consistent, rather than globally.

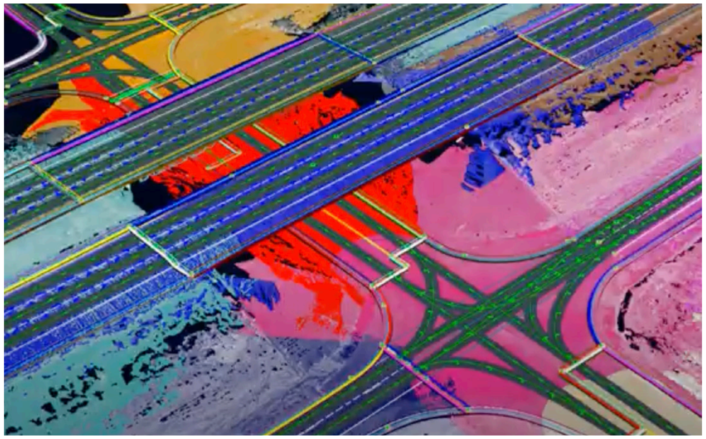

The Aurora Atlas is composed of individual tiles that are easy to update and maintain.

In globally consistent maps, all data is laid out in relation to a single frame, whereas locally consistent maps such the Aurora Atlas, are laid out in relation to multiple frames. Using locally consistent maps allows us to update the Aurora Atlas on a granular scale, without impacting the map as a whole. For example, we can update a change to a driving lane by only changing affected frames, not the entire map. This means that we can update small portions of maps quickly, ensuring that our maps are up-to-date and true to the roads encountered by vehicles traveling along them.

Another key differentiator is that we build all of our Aurora Atlas content entirely in-house to ensure we have the quality of data we need exactly when we need it. This also allows us to more rapidly iterate on the Aurora Atlas to improve the Aurora Driver, and to rapidly build accurate, lightweight maps in new areas and push updates to the on-vehicle Aurora Atlas in near-real time.

Our onboard localization system is responsible for determining the vehicle's position relative to the Aurora Atlas. It 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, dynamic road actors, and environmental changes.

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 driving controls and crash protection necessary for U.S. roads.

We have integrated the Aurora Driver into multiple vehicle platforms that meet or exceed 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. This collaboration allows us to better prepare the Aurora Driver for integration with future vehicle platforms, and helps our partners design their vehicles to directly accept 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.


Aurora

2 Fail Safe

G2
Fail-Safe

The self-driving vehicle is acceptably safe in presence of faults and failures

Fault Management System

As diligently as we are working to build durable, robust, reliable systems, we must anticipate that system failures will occur. Aurora's Fault Management system allows us to continuously detect and safely respond to such faults.

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 and implemented to ensure that fault monitor state changes are always detected whenever a checklist item goes from checked to unchecked, including information on when the fault was detected.

When failures occur, the fault management framework communicates to the vehicle operator so that they can take over if necessary and logs fault diagnostic information. The fault management framework also maps individual item failures to requirements-driven, predetermined fault responses, initiating the appropriate fault response and preventing autonomy engagement when appropriate.

Today, we handle faults by alerting vehicle operators so that they can resume control of the vehicle in the presence of a failure. We ensure that the vehicle operators are able to do this by providing them robust training on system limitations and alerting them via audio and visual cues when a fault is detected or when the Aurora Driver stops sending data to the vehicle.

But our ultimate goal is a vehicle that does not depend on a vehicle operator to manage such a situation. Our solution is to develop the Aurora Driver's Fault Management System (FMS) such that the vehicle can detect faults and operate safely when something goes wrong. It achieves this by performing a minimal risk maneuver (MRM)[5] such as pulling over to achieve a minimal risk condition (MRC)[6] by being outside the flow of traffic.

In the event of a loss of a function or component, the Aurora Driver is capable of reaching a MRC by executing MRMs because fault tolerance and the ability to fail safely are built into the software, hardware, and embed-

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.

ded 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.

Redundancy

It is common practice to use separate, isolated, and 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 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.

The Aurora Driver can only experience a complete failure if the two separate and independent failures occurred roughly at the same time. Aurora designs its safety critical systems with appropriate redundancy in order to provide functionality even after a fault or failure, at which point the FMS will take over.

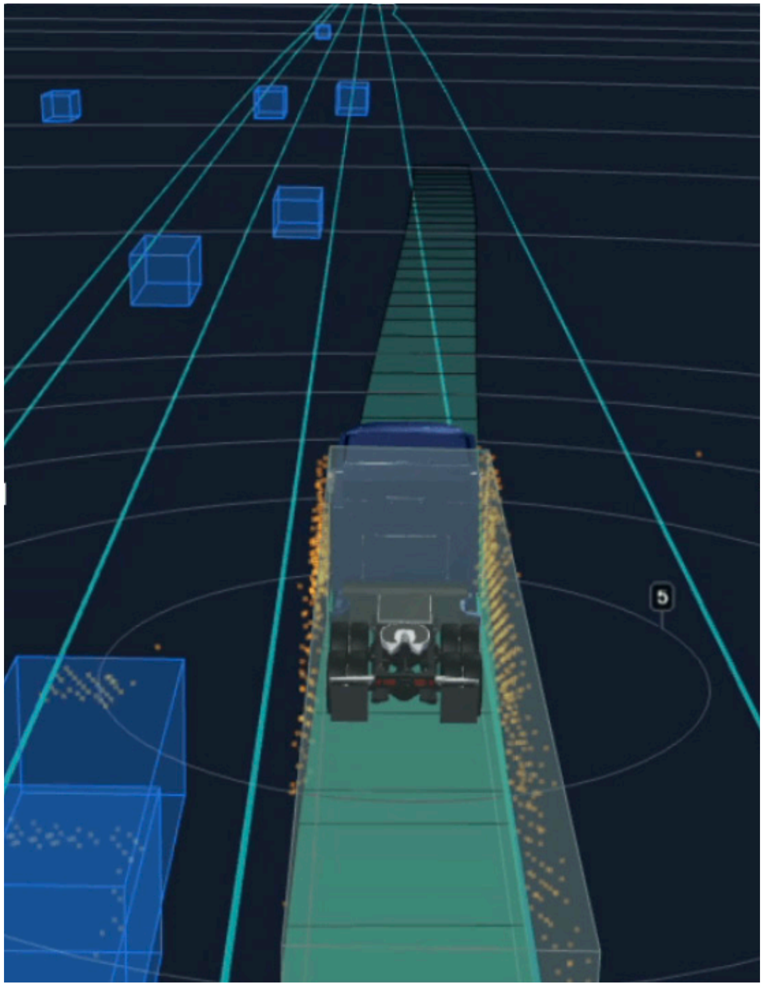

Aurora's Fault Management System will enable our autonomous vehicles to pull over to the shoulder when a problem occurs on the road, like a damaged sensor.

Fault Detection

As mentioned above, Aurora's FMS actively monitors the health of the vehicle, including the self-driving software, sensors, and onboard computer. Each component of the Aurora Driver continuously reports diagnostic health checks to the other components, ensuring that all systems are meeting the right conditions for safe autonomous operation. These health checks correspond to preconditions for nominal execution, and a failed health check will trigger a fault detection. 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 normal operations, the FMS will initiate a mitigation strategy.

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 MRC and then the vehicle will execute that strategy according to the vehicle's environment. Once the MRM has been executed, the vehicle will have achieved the MRC. For example, depending on the situation, the MRC could be the vehicle on the roadway shoulder with the vehicle immobilized and hazard lights flashing.


Aurora

3 Continuously Improving

G3

Continuously Improving

All identified potential safety issues posing an unreasonable risk to safety are evaluated, and resolved iwht appropriate corrective and preventative actions

Safety Performance Indicators

Safety performance indicators are a subset of the broad category of safety metrics that provide meaningful signals on safety risk and inform the improvement of respective process or performance.

Safety performance indicators are a subset of the broad category of safety metrics that provide meaningful signals on safety risk and inform the improvement of respective process or performance. This particular type of safety metric is mapped to one or many safety case claims, and used to monitor for possible safety process and safety performance deficiencies.

At Aurora, we continue to develop safety metrics that could become safety performance indicators to understand not only the performance of the Aurora Driver while in autonomy, but also our operations and the organization itself for driverless operations. As the value provided by each metric can vary at different points of the development process as the performance of the system improves and as capabilities are implemented, we purposefully monitor and analyze multiple metrics throughout the entire development and testing process to better understand our development efforts and identify potential areas for continual focus and improvement.

Our approach to our broad category of safety metrics was informed by key industry standards (see the "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 performance targets are computed using data collected from all of our testing activities (simulation, track, and road).

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

- Scenario Test Results
- Execution of Behavioral Competencies
- Vehicle Dynamics (Velocity, Jerk, Acceleration)
- Policy & Reactionary Disengagements[7]
- Spatial Separation from Other Road Users
- System Faults & Responses
- System Reaction Times

7. Policy disengagements are when our vehicle operators proactively disengage the Aurora Driver when they face an on-road situation that is out of scope for the software in development — that is, any situation that the Aurora Driver has not yet been taught to handle. Reactionary disengagements are situations when our vehicle operators disengaged the Aurora Driver because they believed there was a chance that an unsafe situation might occur.

Cybersecurity

Cybersecurity risks are constantly evolving, so continuous improvement in handling them is critical. By proactively exploring risks, investing in solutions, and collaborating with our industry partners, we regularly incorporate security upgrades across our vehicle platforms in order to harden them against threats. We are dedicated to advancing security approaches and capabilities within all components to improve the security posture for future self-driving vehicles across the industry. Later in the Resilient section of this report, we discuss cybersecurity and its role in developing the Aurora Driver in more detail.

Safety Concern Reporting

We strive to continuously improve our Safety Concern Reporting program because of its importance to our overall Safety Risk Management[8] program and because, in a rapidly developing product, new risks and concerns emerge over time. Our Safety Concern Reporting program provides a clear and easily accessible mechanism for any employee, contractor, or vendor at Aurora, regardless of level or tenure, to report potential safety issues to our safety team. Team members can also report their concerns to their lead, manager, or supervisor on duty, encouraging open and transparent communication throughout the organization. Safety Concern Reports are then triaged, tracked, and monitored until resolution. We will discuss this program in more detail later in this report.

8. Safety Risk Management is a major component of Aurora's Safety Management System (discussed in greater detail in the Trustworthy section below) and includes the identification, assessment, and mitigation of operational safety risks.


Aurora
CERTIFIED
CLEAN IDLE

4 Resilient

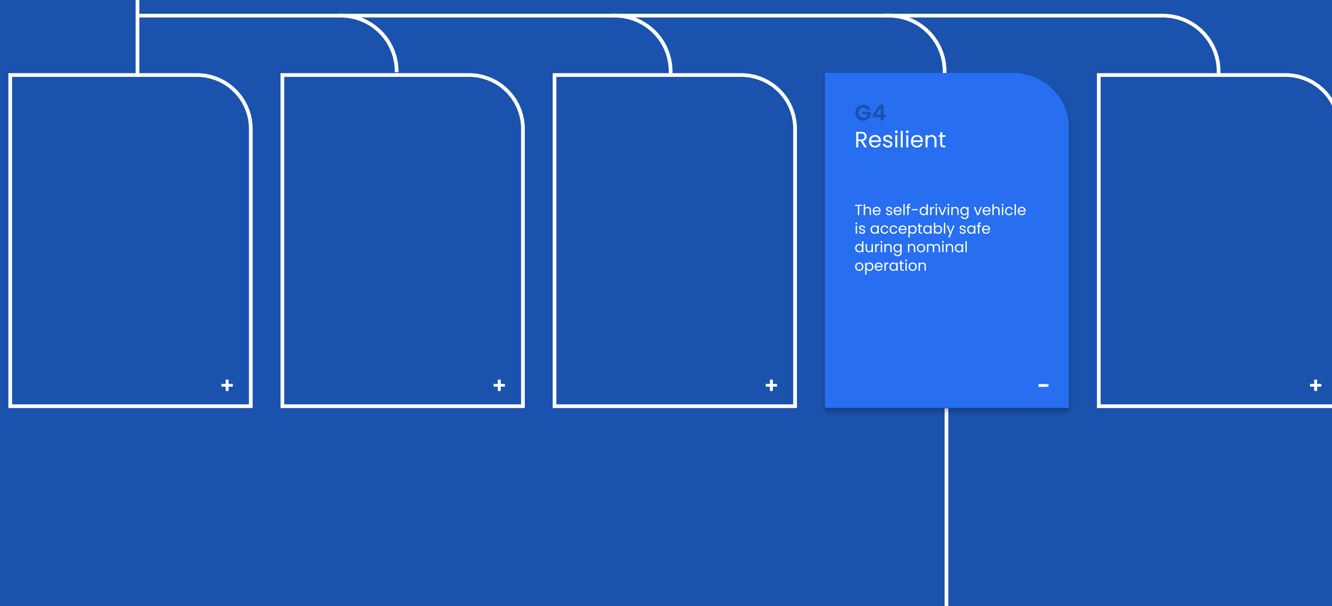
G4
Resilient
The self-driving vehicle is acceptably safe during nominal operation

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.[9]

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 principals, Aurora has developed an extensive and adaptive security approach, aligned with best practices and standards, to secure the extremely 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, a potential vehicle safety concern.

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 narratives that are addressed cross functionally with our partners and across the company.

These narratives, and the controls they derive, serve as a blueprint for the components that must be assessed along with the relative depth for each. Inspired by guidance from NIST,[10] NHTSA, and industry groups, this approach enables Aurora to address security from both a product and process perspective, as well as providing defense in depth through layered controls.

9. Ross, Ron, Mark Winstead, Michael McEvilley. Engineering Trustworthy Secure Systems. NIST: SP 800-160 Vol. 1 Rev. 1, November, 2022.

10. 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.

Trust the Operation of the Aurora Driver

The Trust the Operation of the Aurora Driver theme is composed of four narratives—

1. **Build, Deploy, and Activate** encompasses how we build, test, release, deploy, and activate software for the Aurora Driver platform. This narrative 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. Within the onboard systems, Aurora's systems employ concepts such as trusted software and firmware activation where updates are cryptographically validated.
2. **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.
3. **Engage Autonomy** encompasses the various processes and checks needed to activate and engage the Aurora Driver. Currently, our vehicles are piloted by human vehicle operators that must press an in-cabin button to engage autonomy. In the future, this will transition to a trusted software process controlled by the Aurora Driver onboard and offboard systems.
4. **Trusted Offboard Actions** refers to secure communication between remote specialists and services, and the Aurora Driver. 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 specialists, and vehicle operators also serve as an important line of defense in mitigating security threats. These professionals are trained to detect, annotate, and diagnose any potential irregularities in system and vehicle performance.

Detect, Respond, and Recover

The strategic security risk-reducing architectural changes theme includes two foundational narratives:

1. **Identifying privileged access** to systems or resources is a central security best practice and this narrative focuses on identifying the necessary authorization needed to support the current needs of Aurora Driver development as well as longer term operational needs.
2. **Security detection and response** focuses on ensuring that a security incident can be detected (either through onboard or offboard means), 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 the Fail-Safe section on fault management for a more detailed description of how the Aurora Driver achieves a minimal risk condition.)

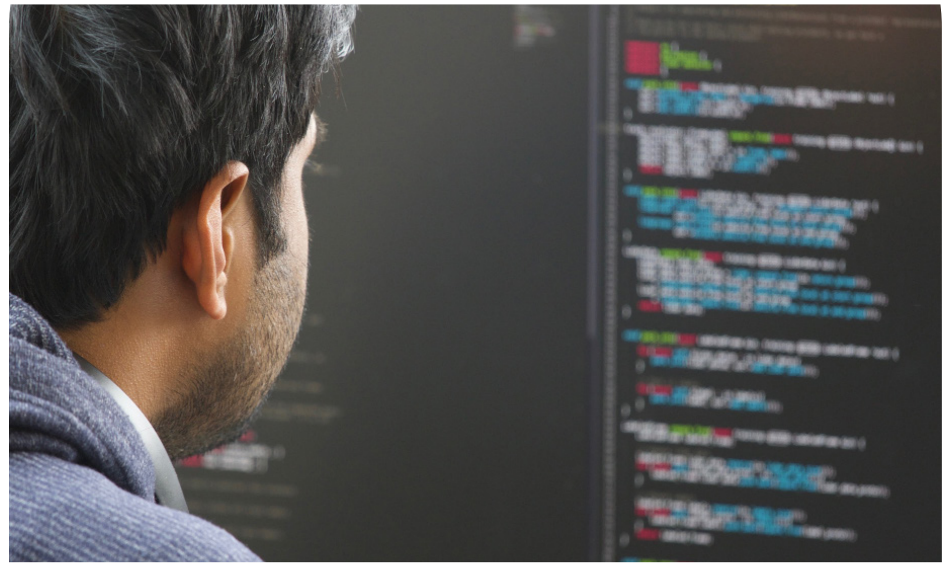

Incident Response

In the event of an incident, Aurora will initiate a series of post-crash procedures which include supporting our vehicle operators, communicating with first responders, and deploying representatives from our response teams to the location. We are respectful of and reliant on the important role played by federal officials, public safety officials, including law enforcement, fire and rescue officials, and emergency medical technicians, in incident response. We proactively engage with public safety officials in the locations where we are operating to ensure that our operations and incident response procedures are complementary to those of public safety agencies.

Our teams follow a defined checklist of actions to ensure on-scene safety as well as provide post-incident support. In the event of an incident our vehicle operators are trained to disengage the Aurora Driver, and if medical attention is needed, to call 911. The vehicle operator will maneuver the vehicle to a safe location away from moving traffic if possible, activate hazard lights, and notify Aurora dispatch, who may assist with data requests, towing, or related post-incident needs.

We have designed the Aurora Driver so that first responders who arrive at the scene can confirm whether the Aurora Driver's engaged by referring to a light bar on the center instrument panel to determine the vehicle state, and disengage as necessary by using an intuitive Aurora Driver shutoff button. This will ensure the vehicle cannot enter self-driving mode. Additional information is available in our Public Safety Official and First Responder Interaction Plans.[11]

Data Recording

The Aurora Driver features a data-logging system which stores raw sensor information as well as other vehicular data, including the operating state of the self-driving system. The Aurora Driver also logs vehicle performance, functionality of the sensor suite, and other data. 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 predefined data from the vehicle that allows us to reconstruct the event, as well as provide the data to the relevant government agencies, when required.

11. Aurora's Public Safety Official and First Responder Interaction Plans (June 2022): Peterbilt 579 and Toyota Sienna Hybrid.


Aurora

5 Trustworthy

G5
Trustworthy

The self-driving enterprise is trustworthy

The Trustworthy principle focuses on ensuring that Aurora is a responsible company that embodies safety. Building safety into the core of our self-driving technology is essential, but ultimately we need to earn the public's trust to deploy it broadly.

This is the foundation of our Trustworthy principle, which guides how our safety culture must support dependable and responsible autonomous vehicle development.

We are committed to a systematic approach to managing safety that will, among other benefits, allow us to further build upon our holistic safety culture. Aurora's safety culture shapes who we hire, how we work, and how we develop our products. We based this robust organizational approach to safety management on best practices found in other safety-critical industries, like aviation, nuclear, and rail, and then tailored it for our culture, operations, and approach to product development.

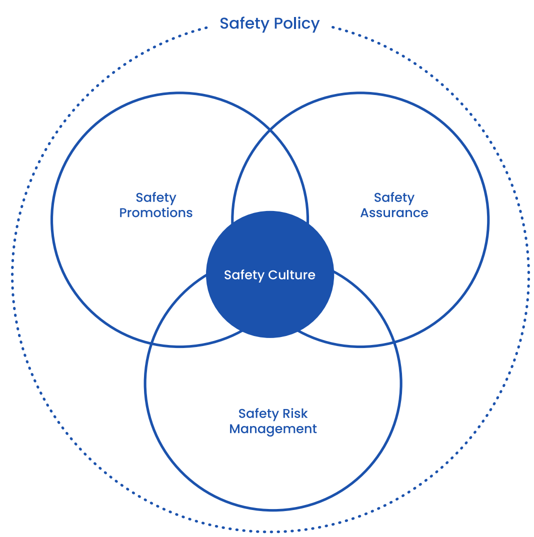

Safety of the Organization / Safety Management System

The Aurora Safety Management System (SMS) allows us to address safety performance by proactively managing safety risk. We do this by presenting relevant safety information to the right person at the right time and ensuring accountability and transparency for every known safety issue across the company.

Our SMS includes mechanisms for all employees to report safety concerns (our Safety Concern Reporting System),[12] a dedicated governance committee (our Safety Review Board),[13] risk assessment processes, safety issue management, a varied and dynamic safety communications program, an independent safety organization, annual safety culture surveys, formal safety education, and a variety of assurance programs.

In addition to our SMS, we have also adopted other tried and tested safety programs, such as ones dedicated to deep collaboration across the company. For example, the Operational Safety Program, an arm of the Safety department that specializes in human factors, fatigue management, and incident response, is fully integrated within our Operations department to foster safe on-road testing and identify opportunities for growth and continuous improvement. More broadly, our Safety Investigation Program works with several teams to root cause reported or observed issues, and facilitate mitigations when necessary.

While these programs stem from Aurora's independent Safety department, a dedicated Safety Culture program encourages all employees to collaborate and participate in safety programs through frequent communications, training, and engagement. This shared responsibility is exemplified in our grounding policy, where anyone in the company can request to stop our on-road operations if they believe there is unmitigated fleet-wide risk exposure. Concerns or system-issue discoveries are elevated to a designated cross-functional response team to review and make the necessary time-critical and local-context-rich decision on grounding needs.

Similarly, we educate and empower all employees to submit Safety Concerns whenever they have a concern, or even a suggestion, on safety at Aurora without fear of reprisal. We are incredibly proud of the participation we have seen in Safety Concern submission and believe it to be a clear indicator of a strong and positive safety culture.

Finally, as part of our commitment to Trustworthiness, we frequently engage with stakeholders outside of Aurora, including federal and state regulators, law enforcement agencies (e.g., NHTSA, the Federal Motor Carrier Safety Administration (FMCSA), National Transportation Safety Board (NTSB), Texas Department of Transportation (TxDOT), Texas Department of Public Safety (TxDPS), and the Commercial Vehicle Safety Alliance (CVSA)), and

12. Our Safety Concern Reporting System provides a clear and easily accessible mechanism for reporting potential safety issues to our safety team, and can be used by any employee, contractor, or vendor at Aurora, regardless of level or tenure. Team members can also report their concerns to their lead, manager, or supervisor on duty, encouraging open and transparent communication throughout the organization.
13. Our Safety Review Board (SRB) is chaired by our CEO, who has been designated as the Accountable Executive for safety. The SRB provides the highest level of safety risk deliberation, decision making, and safety performance oversight for the Company. While managers and directors are empowered to make decisions surrounding safety risk to certain identified risk levels, the SRB reviews the high level risks to determine whether proposed actions are acceptable or whether additional controls are necessary to bring the risk to an acceptable level. The SRB meets regularly to review data-driven metrics on safety performance.

select safety advocacy groups. Through these regular engagements, we are not only able to keep these external stakeholders apprised of our progress, but also consider input they may provide and work with them to develop new policies and procedures.

In 2021, we established an external Safety Advisory Board of esteemed transportation leaders and have been encouraged by their feedback and guidance.[14] We also participate actively in multiple voluntary standards development groups like the Automated Vehicle Safety Consortium (AVSC), SAE, the Institute of Electrical and Electronics Engineers (IEEE), and Underwriters Laboratories (UL).

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

Safety Metrics Reporting

Safety metrics and safety performance metrics serve additional purposes beyond Aurora's own product development and performance monitoring. Aurora reports certain data, such as collision information or other performance anomalies, directly to Federal and State regulators to satisfy mandated data collection requirements[15] or to support granted exemptions from Federal standards, when applicable. These data help government stakeholders learn about the real-world performance of self-driving vehicles and inform their regulatory responsibilities.

14. Learn more about our Safety Advisory Board in this blog post
15. See, e.g., NHTSA First Amended Standing General Order 2021-01 (Aug. 5, 2021) or 13 Cal. Code Regs. § 227.48 (Reporting Collisions) and § 227.50 (Reporting Disengagement of Autonomous Mode).

Industry Best Practices and Standards

Safety metrics and safety performance metrics serve additional purposes beyond Aurora's own product development and performance monitoring. Aurora reports certain data, such as collision information or other performance anomalies, directly to Federal and State regulators to satisfy mandated data collection requirements[15] or to support granted exemptions from Federal standards, when applicable. These data help government stakeholders learn about the real-world performance of self-driving vehicles and inform their regulatory responsibilities.

Aurora believes that being a leader in the development of autonomous vehicles includes collaborating with others in the industry. A key part of our approach to safety is making components of our safety work available to industry stakeholders for their own adaptation, as we did with the public release of our Safety Case Framework[16] It also means sharing 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 sharing safety know-how and collaborating on standards development. These activities are important because we expect collaborative standards development to help shape how self-driving systems are safely developed and deployed. Given the complexity of the challenge the autonomous vehicle industry is tackling, we don't think there will be a single standard that addresses all aspects of safety for self-driving vehicles, so we've taken an active role in AVSC, SAE, IEEE, and Underwriters Laboratories. We anticipate that, much like in the past, voluntary industry standards for self-driving technologies will eventually inform the development of specific federal guidance and/or FMVSS governing those technologies. We support NHTSA's current role in the regulation of motor vehicle safety and equipment and its efforts to develop a technology-neutral approach to guiding and regulating self-driving technology.

The U.S. Department of Transportation Automated Vehicles Comprehensive Plan[17] identified a set of twenty automotive best practices and standards which individually address different components or processes for a self-driving system or enterprise. Further, 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. In addition, standards from other industries have utility in the development and deployment of self-driving vehicles. Though many of these standards were not written specifically with automated driving in mind, Aurora is nonetheless committed to reviewing, interpreting, and applying relevant voluntary industry standards and best practices across our organization.

Safety is an ongoing and iterative process, and as a result,

16. Learn more about our Safety Case Framework in this blog post

17. U.S. DOT, 2021

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 ensure the vehicle is operating safely, whether 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. They close the feedback loop for our developers by providing them 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.

Vehicle Operator Roles in Autonomous Operations

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

The pilot 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 making contact with the steering wheel or pressing on the throttle or brake pedals which disengages autonomy.

The co-pilot 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 each mission is executed, the co-pilot monitors the performance of the Aurora Driver to 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.

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 a weekly refresher training program, which includes material on any new process or procedures from the past week.

All pilots for truck operations possess a valid Commercial Driver's License and have years of experience as a Class A Commercial Driver — most have more than 10 years of experience. All new pilots are enrolled in new hire operator training, as well as continuous education.

Our new hire training currently consists of three components — live classroom instruction, in-vehicle training, and evaluations and exams. Live classroom training leverages instructor presentations, independent reading, video demonstrations, and quizzes and exams to ensure trainee proficiency. In-vehicle training brings classroom learning together with coaching where instructors provide in-the-moment feedback. Trainees are evaluated and certified at each stage before they are allowed to move on to the next level of training. New hires must score 80% or higher on all quizzes and written evaluations.

Our training program features a combination of class-room learning to understand the core concepts, including the operational limitations and how to disengage the Aurora Driver, ODD operating parameters, vehicle controls and safety components of the vehicle platform, pre- and post-trip procedures, and trip planning, among other concepts. They also gain hands-on experience inside the vehicle on a closed course track. Our in-vehicle training includes a driver safety training program with the goal of reducing the potential for a collision.

Finally, our new pilots must demonstrate proficiency in our Vehicle Controllability Clinic. The purpose of the clinic is to highlight best practices, methods, and techniques that maximize safe, smooth, and compliant driving. This training is designed to help our pilots optimize their ability to react, understand how to reduce risk factors, maximize time to respond to changing traffic patterns, and understand how vehicle operation impacts passenger safety.

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:** Our truck pilots adhere to the FMCSA's Hours of Service regulation. 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 their mobile devices and/or smartwatch 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.

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 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, these 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 mission or route types. Therefore, a detailed FRMP aids in the development of risk controls to manage individual and task-related fatigue.

Light Bar Visual Indicators

White
Indicates not ready
(manual control)

Blue
Indicates ready state
(manual control)

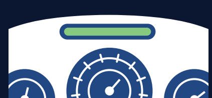

Green
Indicates healthy state
(self-driving engaged)

Red
Indicates unhealthy state
manual takeover required
(self-driving engaged)

Human-machine Interface

Human drivers are constantly receiving new information from the driving environment, processing this information, and making informed decisions. Audio-visual alerts within the vehicle are designed to provide vehicle operators with a quick and easy understanding of the state of the self-driving system. Additionally, the audio and visual indicators improve vehicle operator response time.

Visual Safety Features

Visual alerts with critical information about the state of the Aurora Driver are visible to both the pilot and co-pilot via a light bar on the center instrument panel. A solid green light indicates that the self-driving system is engaged and actively operating in self-driving. When the self-driving system experiences a fault, the light bar automatically turns red.

Audio Safety Features

Unambiguous audio indications are played for certain state transitions and when the self-driving system is faulted. Examples of when the audio is played include:

- When takeover is required;
- When engagement is attempted but fails;
- When the system disengages, regardless of cause;
- When the system engages.

Vehicle Equipment Inspection

Each day, before and after each trip, we inspect vehicles for any internal or external vehicle issues associated with the base vehicle or the self-driving system. We inspect each vehicle's parts and accessories, including brakes, windshield wipers, steering, lighting devices and reflectors, tires, rear vision mirrors, and emergency equipment, and the vehicle's sensor suite and computers. The inspectors will note any issue which could affect the safe operation of the vehicle, and our maintenance reporting system prohibits operation of the vehicle while such conditions persist. Our service technicians also conduct additional monthly routine maintenance inspections of internal and external vehicle equipment and systems.

Developing the Aurora Driver is a complex endeavor. 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 — proficient, fail-safe, continuously improving, resilient, and trustworthy. It is through these principles that we are focusing our efforts to deliver the benefits of self-driving technology safely, quickly, and broadly.


Aurora

Appendix

Appendix

Aurora VSSA Contents	NHTSA VSSA elements
1 Executive Summary	
2 Introduction	
3 Proficient	
a. Aurora's Operational Design Domain	ODD
b. The Aurora Driver	System safety, OEDR
c. Testing the Aurora Driver	System safety, validation methods
d. Vehicle control limits	System Safety
e. HD mapping and localization	ODD
f. Safety of the vehicle platform	Crashworthiness
4 Fail Safe	Fallback, Post Crash ADS Behavior
a. Fault Management System	Fallback
5 Continuously Improving	
a. Safety Performance Indicators	Validation methods
b. Cybersecurity	Cybersecurity
c. Safety concern reporting	
6 Resilient	
a. Cybersecurity	Cybersecurity
b. Incident response	Data Recording, Post Crash ADS Behavior
7 Trustworthy	
a. Safety of the organization / SMS	Consumer education and training
b. Safety metrics reporting	Federal, State, and Local laws
c. Industry best practices and standards	
d. Vehicle operators	HMI, Fallback

Acronyms

ADS
Automated Driving System

AM
Amplitude Modulation

AVSC
Automated Vehicle Safety Consortium

CVSA
Commercial Vehicle Safety Alliance

NHTSA
National Highway Traffic Safety Administration

NTSB
National Transportation Safety Board

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

GPS
Global Positioning System

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

TxDOT
Texas Department of Transportation

TxDPS
Texas Department of Public Safety

UL
Underwriters Laboratories

VSSA
Voluntary Safety Self-Assessment


Aurora