

NVIDIA Autonomous Vehicles Safety Report



Our Mission

The next generation of transportation is autonomous. From shared and personal vehicles, to long- and short-distance travel, to delivery and logistics, autonomy will fundamentally improve the way the world moves. At NVIDIA, our automotive team's mission is to develop self-driving technology that enables safer, less-congested roads, and mobility for all.

“Safety is the most important aspect of a self-driving vehicle. NVIDIA's creation of a safe self-driving platform is one of our greatest endeavors, and provides a critical ingredient for automakers to bring autonomous vehicles to market.”

Jensen Huang, NVIDIA founder and CEO



Contents

Introduction	4	Architected for Safety	17	Summary	24
Level 4 Driving and the Breakthroughs Unlocking It	6	Hardware	19	Appendices	25
Pillars of Safe Autonomous Driving	9	Software	19	NVIDIA Activity in Expert Groups	
Artificial Intelligence Design and Implementation Platform	9	Vehicles and Sensors	20	National and International Safety Regulations and Recommendations	
Development Infrastructure for Deep Learning	12	Data Center	21	NHTSA Elements of Safety	
Physically Accurate Sensor Simulation for AV Development	13	On-Road Testing	22	References	
Best-In-Class Pervasive Safety and Cybersecurity Program	14	Developer Training and Education	23		

Introduction

NVIDIA pioneered accelerated computing to tackle challenges no one else can solve. Our work in AI and industrial digitalization is profoundly impacting society and transforming the world's largest industries—from gaming to robotics, to life-saving healthcare and climate change, to virtual worlds where we can all connect and create.

NVIDIA is also applying our technology-driven vision, computational performance, and energy efficiency to the transportation industry—helping vehicle makers around the world realize the dream of safe and reliable autonomous vehicles. From concept and design, to engineering and production, to sales and service, NVIDIA technologies are streamlining the entire automotive industry's workflow.

In particular, automated and autonomous vehicles are poised to transform the transportation industry. They have the potential to dramatically reduce injuries and fatalities from collisions, alleviate traffic congestion, increase productivity, and provide mobility to those who are unable to drive.

Breakthroughs in AI and accelerated computing are opening future fleets to dramatic new functionality,



fundamentally transforming the vehicle architecture for the first time in decades into a truly AI-defined architecture. Like all modern computing devices, these intelligent vehicles are supported by a large team of AI experts and software engineers, dedicated to improving the performance and capability of the car as technology advances. Capabilities and services can be added using over-the-air updates throughout the entire life of the car.

NVIDIA works with vehicle makers, suppliers, sensor manufacturers, and startups around the world. We provide the systems architecture, AI supercomputing hardware, and full software stack required to build all types of vehicles—from AI-assisted cars and trucks to fully autonomous shuttles and robotaxis. With an open, modular architecture that spans from

the cloud to the car, manufacturers can use select solutions or the entire development pipeline.

It all starts with NVIDIA DRIVE®, our highly scalable platform that can enable all levels of autonomous driving as defined by the Society of Automotive Engineers (SAE). These range from advanced driver-assistance system features (SAE Level 2: driver-assisted) through robotaxis (SAE Level 5: full automation).

The computational requirements for fully autonomous driving are enormous—easily up to 100 times higher than advanced vehicles in production today. With NVIDIA DRIVE, our partners can achieve the highest levels of safety with an architecture featuring diversity and redundancy in the computing hardware, sensor suites, and software stack.

To streamline development, we've created a single software-defined scalable architecture that advances each level of autonomy with additional hardware and software while preserving the core architecture. The same strategy applies for safety. With additional modular hardware and software, the level of safety achieved scales to match the more rigorous requirements of advanced levels of autonomy.

NVIDIA has created essential technologies for building robust systems for the research, development, and deployment of self-driving vehicles, spanning from the data center to the car. We offer a range of hardware and software solutions, from powerful GPUs and servers to a complete AI

training infrastructure and in-vehicle autonomous driving supercomputers. We also support academic research and early-stage developers, partnering with dozens of universities worldwide and teaching courses on AI development at our Deep Learning Institute (DLI). As we identify challenges, we turn them into opportunities and build solutions.


Introducing Halos

In March 2025, NVIDIA announced Halos, a system that unifies the company's industry-leading lineup of automotive hardware and software technologies and AV safety research under one identity. This provides a truly comprehensive safety system for the automotive ecosystem.

The Halos system combines autonomous vehicle hardware, software, and design principles to help ensure safety, from the cloud to the car. It encompasses all NVIDIA DRIVE automotive products and technologies, as well as an advanced algorithmic safety layer. Halos also includes the NVIDIA AI Systems Inspection Lab to help automotive ecosystem partners navigate evolving industry standards for autonomous vehicle safety. The lab inspects and verifies that automotive partner software and systems on the NVIDIA DRIVE AGX™ platform meet the automotive industry's stringent safety and cybersecurity standards, including AI functional safety.

The lab has been accredited by the ANSI National Accreditation Board (ANAB) according to the ISO/IEC 17020 assessment for standards, including:

- > Functional safety (ISO 26262)
- > SOTIF (ISO 21448)
- > Cybersecurity (ISO 21434)
- > UN-R regulations, including UN-R 79, UN-R 13-H, UN-R 152, UN-R 155, UN-R 157, and UN-R 171
- > AI functional safety (ISO PAS 8800 and ISO/IEC TR 5469)




Level 4 Driving and the Breakthroughs Unlocking It

When the Society of Automotive Engineers established its [framework for vehicle autonomy](#) in 2014, it created the industry-standard roadmap for self-driving technology.

The levels of automation progress from level 1 (driver assistance) to level 2 (partial automation), level 3 (conditional automation), level 4 (high automation), and level 5 (full automation).

Predicting when each level would arrive proved more challenging than defining them. This uncertainty created industry-wide anticipation, as breakthroughs seemed perpetually just around the corner.

That dynamic has shifted dramatically in recent years, with more progress in autonomous driving made in the past three to four years than in the previous decade. Continue reading to learn about the advancements that have made such rapid progress possible.



What Is Level 4 Autonomous Driving?

Level 4 autonomous driving enables vehicles to handle all driving tasks within specific operating zones—such as certain cities or routes—without the need for human intervention. This high-automation level uses AI breakthroughs, including foundation models, end-to-end architectures, and reasoning models to navigate complex scenarios.

Today, level 4 “high automation” is bringing the vision of [autonomous driving](#) closer to a scalable, commercially viable reality.

How AI Is Advancing Autonomous Vehicles

Six major AI breakthroughs are converging to accelerate level 4 autonomy:

1. Foundation Models

[Foundation models](#) can tap internet-scale knowledge, not just proprietary driving fleet data.

When humans learn to drive at, say, 18 years old, they’re bringing 18 years of world experience to the endeavor. Similarly, foundation models bring a breadth of general world knowledge that informs their ability to understand unusual scenarios and predict outcomes.

With foundation models, a vehicle can reason its way through scenarios it has never seen before—such as encountering a mattress in the road or a ball rolling into the street—drawing on information learned from vast training datasets.

2. End-to-End Architectures

Traditional autonomous driving systems used separate modules for perception, planning and control, losing information at each handoff.

[End-to-end autonomy](#) architectures have the potential to change that. With end-to-end architectures, a single network processes sensor inputs directly into driving decisions, maintaining context throughout. While the concept of end-to-end architectures isn’t new, architectural advancements and improved training methodologies are finally making this paradigm viable, resulting in better autonomous decision-making with less engineering complexity.

3. Reasoning Models

[Reasoning vision-language-action \(VLA\)](#) models integrate diverse perceptual inputs, language understanding, and action generation with step-by-step reasoning. This enables them to break down complex situations, evaluate multiple possible outcomes, and decide on the best course of action—much like humans do.





Systems powered by reasoning models deliver far greater reliability and performance, with explainable, step-by-step decision-making. For autonomous vehicles, this means the ability to flag unusual decision patterns for real-time safety monitoring, as well as post-incident debugging to reveal why a vehicle took a particular action. This improves the performance of autonomous vehicles while building user trust.

4. Simulation

With physical testing alone, it would take decades to test a driving policy in every possible driving scenario, if achievable at all. Enter [simulation](#).

Technologies like [neural reconstruction](#) can be used to create interactive simulations from real-world sensor data, while world models like [NVIDIA Cosmos™ Predict and Transfer](#) produce unlimited novel situations for training and testing [autonomous vehicles](#).

With these technologies, developers can use text prompts to generate new weather and road conditions, change lighting, or introduce obstacles to simulate new scenarios and test driving policies in novel conditions.

5. Compute Power

None of these advances would be possible without sufficient computational power. The [NVIDIA DRIVE AGX](#) and [NVIDIA DGX™](#) platforms have evolved through multiple generations, each designed for today's AI workloads as well as those anticipated years down the road.

Co-optimization matters. Technology must be designed anticipating the computational demands of next-generation AI systems.

6. AI Safety

Safety is foundational for level 4 autonomy, where reliability is the defining characteristic distinguishing it from lower autonomy levels. Recent advances in physical AI safety enable the trustworthy deployment of AI-based autonomy stacks by introducing safety guardrails at the stages of design, deployment, and validation.

For example, NVIDIA's safety architecture guardrails the end-to-end driving model with checks supported by a diverse modular stack, and validation is greatly accelerated by the latest advancements in neural reconstruction.

Why It Matters: Saving Lives and Resources

The stakes extend far beyond technological achievement. Improving vehicle safety can help save lives and conserve significant amounts of money and resources. Level 4 autonomy systematically removes human error, the cause of the vast majority of crashes.

NVIDIA, as a full-stack autonomous vehicle company—from cloud to car—is enabling the broader automotive ecosystem to achieve level 4 autonomy, building on the foundation of its level 2+ stack already in production. In particular, NVIDIA is the only company that offers an end-to-end compute stack for autonomous driving.

The three AI compute platforms critical for autonomy are:

- [NVIDIA DGX](#) for training AI in data centers
- [NVIDIA Omniverse™](#) and [Cosmos](#) to simulate worlds and generate synthetic data for testing and validation
- [NVIDIA DRIVE AGX](#) in-vehicle computers to process sensor data in real time

Together, these platforms form a feedback loop for learning, testing, and deployment that tightens the cycle of innovation while keeping safety front and center.

Pillars of Safe Autonomous Driving

NVIDIA offers a unified hardware and software architecture throughout its autonomous vehicle research, design, and deployment infrastructure. We deliver the technology to address the four major pillars essential to making safe self-driving vehicles a reality.

- > Artificial Intelligence Design and Implementation Platform
- > Development Infrastructure to Support Deep Learning

- > Physically Accurate Sensor Simulation for AV Development
- > Best-In-Class Pervasive Safety and Cybersecurity Program

Artificial Intelligence Design and Implementation Platform

NVIDIA DRIVE is the world's first scalable AI platform that spans the entire range of autonomous driving, from AI-assisted driving to robotaxis. The platform consists of hardware, software, and firmware that work together to enable the production of automated and self-driving vehicles.


Our platform combines deep learning with traditional software to enable a safe driving experience. With high-performance computing, the vehicle can understand in real time what's happening around it, precisely localize itself, and plan a safe path forward.

Our unified architecture extends from the data center to the vehicle and provides a comprehensive solution addressing the requirements of national and international safety standards.

Deep neural networks (DNNs) are trained on the NVIDIA DGX platform, which incorporates the best of NVIDIA software, infrastructure, and expertise in a modern, unified AI development solution. Then, they're tested and validated in simulation on NVIDIA OVX before being seamlessly deployed to run on our AI computer in the vehicle. NVIDIA OVX is a computing system designed to power large-scale Omniverse digital twins. To operate safely, self-driving vehicles require onboard supercomputers capable of processing all the sensor data in real time.


NVIDIA DRIVE Hardware

Our underlying hardware solutions include:




NVIDIA DRIVE AGX Hyperion

NVIDIA DRIVE AGX Hyperion™ is an end-to-end, modular reference architecture for designing autonomous vehicles. It accelerates development, testing, and validation by integrating DRIVE AI-based compute with a complete sensor suite, which includes exterior and interior cameras, ultrasonics, radars, and lidars.




NVIDIA DRIVE AGX Orin

The NVIDIA DRIVE AGX Orin™ SoC (system-on-a-chip) delivers up to 254 TOPS (trillion operations per second) and is the central computer for intelligent vehicles. It's the ideal solution for powering autonomous driving capabilities, confidence views, digital clusters, and AI cockpits. The DRIVE AGX Orin product family lets developers build, scale, and leverage one development investment across an entire fleet, from Level 2+ systems all the way to Level 5 fully autonomous vehicles.



NVIDIA DRIVE AGX Thor

The DRIVE AGX Thor™ SoC is our next-generation centralized car computer, combining feature-rich cockpit capabilities, plus highly automated and autonomous driving on a single, safe, and secure system. This AV processor uses our latest CPU and GPU advances—including the NVIDIA Blackwell GPU architecture for transformer and generative AI capabilities. DRIVE AGX Thor features 8-bit floating point support (FP8) to deliver an unprecedented 1,000 INT8 TOPS/1,000 FP8 TFLOPS/500 FP16 TFLOPS of performance while reducing overall system cost.



NVIDIA Blackwell Architecture

Powering a new era of computing, the NVIDIA Blackwell platform will enable organizations everywhere to build and run real-time generative AI on trillion-parameter large language models at up to 25X less cost and energy consumption than its predecessor. The Blackwell GPU architecture features transformative technologies for accelerated computing, including the world's most powerful chip.

NVIDIA DRIVE Software Development Kits

Software is what turns a vehicle into an intelligent machine. The open NVIDIA DRIVE SDK provides all the building blocks and algorithmic stacks needed for autonomous driving. This helps developers build and deploy a variety of state-of-the-art AV applications more efficiently, including perception, localization and mapping, planning and control, driver monitoring, and natural language processing.

- ▶ The foundation of the DRIVE software stack is NVIDIA DriveOS™. This is the first safe operating system for in-vehicle accelerated computing with industry-standard safety and security methodologies certified by the globally renowned automotive certification organization, TÜV SÜD. It includes NVIDIA® CUDA® libraries for efficient parallel computing implementations, NVIDIA TensorRT™ for real-time AI inference, and NvMedia for sensor input processing.
- ▶ NVIDIA DriveWorks provides middleware functions on top of NVIDIA DriveOS that are fundamental to autonomous vehicle development. These consist of the sensor abstraction layer (SAL) and sensor plugins, a data recorder, vehicle I/O support, and



a Deep Neural Network (DNN) framework. It's modular, open, and designed to be compliant with automotive industry software standards.

- ▶ The DRIVE SDK enables development of safety-critical perception functions including object detection, classification, tracking, and sensor fusion targeting ASIL-B to ASIL-D integrity levels.
- ▶ NVIDIA offers an AI-assisted driving platform that can handle both highway and urban traffic with the highest levels of safety. It can use the high-performance compute and sensor set of NVIDIA

DRIVE Hyperion to drive from address to address. For those who want to drive, the system also provides active safety features and the ability to intervene in dangerous scenarios.

- ▶ NVIDIA also delivers new, always-on intelligent services for the driver and passengers. The NVIDIA Avatar Cloud Engine (ACE) serves as a digital assistant, making recommendations, helping book reservations, making phone calls, accessing vehicle controls, and providing alerts using natural language.

Development Infrastructure for Deep Learning

In addition to in-vehicle supercomputing hardware, NVIDIA designs and develops the supercomputers used to solve critical challenges faced in the development and deployment of safe AVs. A single test vehicle can generate petabytes of data each year. Capturing, managing, and processing this massive amount of data for an entire fleet of vehicles requires a fundamentally new computing architecture and infrastructure.

NVIDIA AI Training and Simulation

NVIDIA provides the complete data center hardware, software, and workflows needed to develop autonomous driving technology—from raw data collection through validation. It provides the end-to-end building blocks required for neural network development, training and validation, and testing in simulation.

> **NVIDIA DGX Systems:** These are purpose-built AI supercomputers designed to train deep neural networks. They allow for the training of highly complex models needed for autonomous driving on large datasets. DGX systems enable the training of robust AI models capable of handling complex driving scenarios. By training on diverse

and extensive datasets, the models can better generalize to various real-world conditions, enhancing safety.

> **NVIDIA Omniverse Cloud Sensor RTX™:** This set of cloud-based tools for developing autonomous vehicle technologies provides high-fidelity simulation environments with realistic physics and sensor models, greatly enhancing accuracy in testing. Omniverse allows for extensive testing in a virtual environment before real-world deployment. This helps identify and mitigate potential safety issues in a controlled setting, reducing risks during actual operation.

Data Management and Cloud Services

Efficient data management and cloud-based services are critical for autonomous vehicle development:

> **NVIDIA AI Infrastructure:** Tapping into NVIDIA's expertise in high-performance computing and AI, this infrastructure supports the large-scale data processing and storage needs of autonomous vehicle development. NVIDIA AI Infrastructure is a solution for the entire industry. Today, a leading automaker is using over 35,000 GPUs to advance its AV development and testing.



Physically Accurate Sensor Simulation for AV Development


Before any autonomous vehicle can safely navigate on the road, engineers must first train, validate, and test the AI algorithms and other software that enable the vehicle to drive itself. AI-powered autonomous vehicles must be able to respond properly to the incredibly diverse situations they could experience, such as emergency vehicles, pedestrians, animals, and a virtually infinite number of other obstacles—including scenarios that are too dangerous to test in the real world.

In addition, AVs must perform regardless of weather, road, or lighting conditions. But there's no practical way to physically road test vehicles in all these situations, nor is road testing sufficiently controllable, repeatable, exhaustive, or efficient. The ability to test in a realistic simulation environment is essential to providing safe self-driving vehicles. Coupling actual road miles with simulated miles in the data center is the key to developing and validating AVs.

Autonomous vehicle simulation requires extremely tight timing, repeatability, and real-time performance, and it must be able to operate at scale. Additionally, generating data from AV sensor sets in physically based virtual worlds requires tremendous compute loads.

NVIDIA Omniverse Cloud Sensor RTX, built on OpenUSD, is designed to let developers enhance their AV simulation workflows with high-fidelity sensor simulation, physics, and realistic behavior. With these APIs, you can connect to a vast ecosystem of partners building simulation tools for vehicle dynamics and traffic. You can also bring in USD content to expand to new locales and tackle new operational design domains (ODDs).


NVIDIA Omniverse libraries, including Sensor RTX and NuRec, enable physically based and neural rendering of sensors commonly deployed on autonomous vehicles. These include cameras, lidar, radar, and ultrasonic sensors.



Best-In-Class Pervasive Safety and Cybersecurity Program

Safety

Safety is our highest priority at every step of the AV research, development, and deployment process. It begins with a safety methodology that emphasizes diversity and redundancy in the design, validation, verification, and lifetime support of the entire autonomous system. We follow and develop best-in-class solutions in our processes, products, and safety architecture. NVIDIA safety is designed for software-defined autonomy because it accepts, tackles, and takes advantage of the complexity of autonomous vehicles.



To conceptualize our autonomous vehicle safety program, we followed recommendations by the U.S. Department of Transportation's National Highway Traffic Safety Administration in its 2017, 2018, and 2020 publications.¹

Throughout our program, we follow the automotive industry's safety standards from the International Organization for Standardization. These include:

Functional Safety (ISO 26262)

Autonomous vehicles must operate safely when a system fails. For L2/L2+, we must detect and mitigate failures (returning control to the driver), and for L3/L4 we must continue to operate safely and reach a minimal-risk condition. We apply functional safety at all levels of hardware, software, and the system—from the application, through middleware and the operating system, to the board and the chips on the board, to the system providing the autonomous driving functionality.

Safety of the Intended Function (ISO 21448)

A system designed to be functionally safe (ISO 26262)² must also be designed and tested to be performant on all safety-critical metrics related to the intended functionality (ISO 21448)³.

Safety hazards can be present even if the system is functioning as designed, without a malfunction. SOTIF focuses on ensuring the absence of unreasonable risk due to hazards resulting from insufficiencies in the intended functionality or from reasonably foreseeable misuse. For example, perception failures must be deemed sufficiently rare so the autonomous vehicle rarely fails to detect a pedestrian in its path.

Safety and Artificial Intelligence

We actively contribute to ongoing standardization initiatives related to safety of artificial intelligence, such as the ISO PAS 8800⁴ (currently in development), the ISO/IEC TR 5469⁵, and its follow-up ISO/IEC TS 22440⁶ (currently in development).

Federal and International Regulations and Standardization

We adhere to federal and international regulations, including global NCAP (New Car Assessment Program), Euro NCAP, and the United Nations Economic Commission for Europe. We also influence, co-create, and follow standards of the International Standards Organization, the New Vehicle Assessment Program, and SAE, as well as standards from other industries.

We contribute to standardization initiatives of the Institute of Electrical and Electronics Engineers (IEEE), such as IEEE 2846-2022 (on Assumptions for Models in Safety-Related Automated Vehicle Behavior)⁷ and IEEE P2851 (on Exchange/Interoperability Format for Safety Analysis and Safety Verification of IP, SoC, and Mixed Signal ICs⁸).

Beyond complying with federal and industry guidelines, we practice open disclosure and collaboration with industry experts to ensure that we remain up-to-date on all current and future safety issues. We also hold leadership positions in multiple safety working groups to drive the state of the art and explore new research areas, such as safety for AI systems and explainable AI.

Meeting the Highest Standards

To make transportation safer, autonomous vehicles must have processes and underlying systems that meet the highest standards.

Recent significant safety certifications and assessments of NVIDIA automotive products by independent, accredited assessors include the following:

- TÜV SÜD certified the NVIDIA DriveOS 6.0 operating system to be conformant with ISO 26262 Automotive Safety Integrity Level (ASIL) D standards.
- TÜV SÜD granted the ISO 21434 Cybersecurity Process certification to NVIDIA for its automotive system-on-a-chip, platform, and software engineering processes.
- TÜV Rheinland performed an independent United Nations Economic Commission for Europe safety assessment of NVIDIA DRIVE AV related to safety requirements for complex electronic systems.

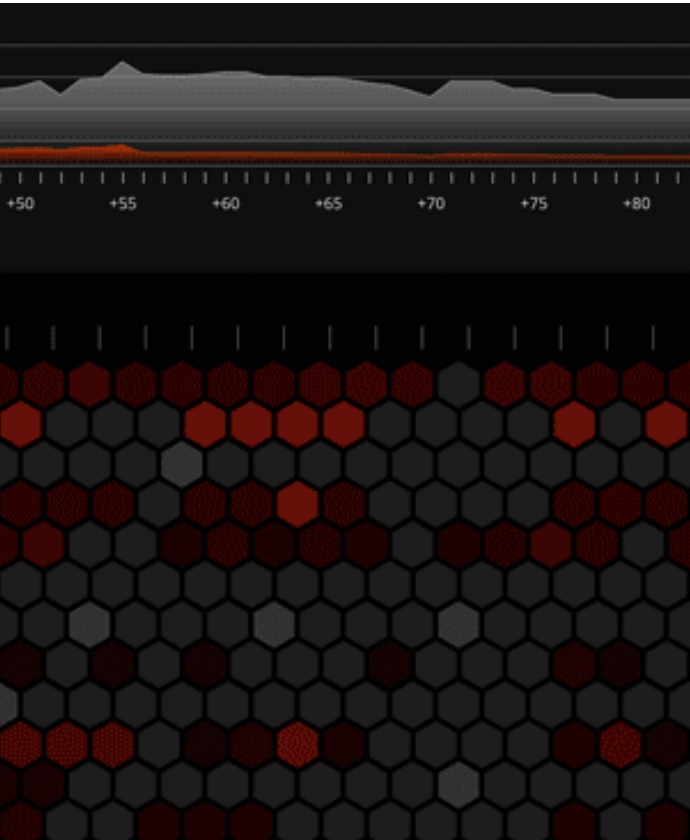
Additional certifications by TÜV SÜD include:

- NVIDIA DRIVE core development processes are certified as ISO 26262 Automotive Safety Integrity Level (ASIL) D compliant.
- The NVIDIA DRIVE AGX Orin SoC completed concept and product assessments and is deemed to meet ISO 26262 ASIL D systematic requirements and ASIL B random fault-management requirements.
- The NVIDIA DRIVE AGX Orin board completed concept assessment and is deemed to meet ISO 26262 ASIL D requirements.
- The NVIDIA DRIVE AGX Orin-based system, which unifies the DRIVE AGX Orin SoC and DRIVE AGX Orin board, completed concept assessment and is deemed to meet ISO 26262 ASIL D requirements.



Cybersecurity

An autonomous vehicle platform can't be considered safe without cybersecurity. Comprehensive security engineering practices and development are essential to deliver on the functional and overall safety required for the automotive industry.



Security breaches can compromise a system's ability to deliver on fundamental safety goals. To deliver a best-in-class automotive security platform with high consumer confidence, we've built a world-class security team and aligned with government and international standards and regulations. We've also built strong partner relationships to remediate security incidents and serve as a good steward in protecting customer data privacy.

NVIDIA follows international and national standards for hardware and software implementations of security functionality, including cryptographic principles. We also adhere to standards set by NIST (National Institute of Standards and Technology) (8) and GDPR (General Data Protection Regulations) (9) to protect the data and privacy of all individuals.

Our cybersecurity team works with the Automotive Information Sharing and Analysis Center (Auto-ISAC), NHTSA, SAE, and the Bureau of Industry and Security (Department of Commerce). We also contribute to the Automatic Identification System (Department of Homeland Security), Federal Information Processing Standards (Federal Information Security Management Act), and Common Criteria standards or specifications.


We follow and maintain a cybersecurity management system as defined in UNECE Regulation No. 155 (10). In addition, we use the ISO/SAE 21434 cybersecurity

process as a guiding principle and align our automotive development practices accordingly. This makes compliance claims easier, besides using processes and practices from other cybersecurity-sensitive industries with ISA/IEC 62443.

We participate in the SAE J3101 standard development, which ensures that the necessary building blocks for cybersecurity are implemented at the hardware and system software levels. We review platform code for security conformance and use static and dynamic code analysis techniques for early detection, and perform penetration and other offensive security techniques for validation. Also, we participate in SAE 8477 to ensure our security testing methodologies evolve over time.

NVIDIA employs a rigorous security development lifecycle into our system design and hazard analysis processes. This includes end-to-end traceability on security requirements, threat models that cover the entire autonomous driving system—hardware, software, manufacturing, and IT infrastructure—ensuring secure design and coding guidelines are in place. The DRIVE platform has multiple layers of defense that provide resiliency against a sustained attack.

The NVIDIA threat intelligence team (NTIP) delivers actionable intelligence to various NVIDIA Business Units (BUs), including automotive, by connecting with



them on their specialized intelligence requirements. At the same time, we delve into a comprehensive plan for dissemination and consumption of finalized intelligence alerts and/or reports.

NVIDIA also maintains a dedicated Product Security Incident Response Team that manages, investigates, and coordinates security vulnerability information internally and with our partners. This allows us to contain and remediate any immediate threats while openly working with our partners to recover from security incidents. We scan all our software for the presence of third-party packages and associated vulnerabilities therein (if any). As part of continuous vulnerability management, development teams are expected to adhere to internal automotive security policies.

To demonstrate that the production, operation, and maintenance of vehicle feature elements are being performed securely in compliance with ISO 21434, security tool management continuously assesses the appropriateness and effectiveness of security controls associated with a software tool. It establishes whether errors associated with the tool can be adequately prevented or detected to minimize adverse impacts to the security of the element.

In addition, we work closely with our suppliers to ensure the components that make up the whole

of an autonomous driving platform provide the necessary security features. Proper cybersecurity of complex platforms becomes assured when all the links in the chain—from raw data to processed input to control actions—meet security requirements. NVIDIA also works with our vendors to ensure they have a cybersecurity reaction capability for new or unfound threats.

Finally, as vehicle systems have a longer in-use lifespan than many other types of computing systems, we use advanced machine learning techniques to detect anomalies in the vehicle's communications and behaviors and provide additional monitoring capabilities for zero-day attacks.

Architected for Safety

Overview

NVIDIA designed the DRIVE AGX platform to ensure that the autonomous vehicle can operate safely within the intended operational design domain (ODD). In situations where the vehicle is outside its defined ODD or conditions dynamically change to

fall outside it, our products enable the vehicle to return to a minimal risk condition (also known as a safe fallback state). For example, if an automated system detects a sudden change, such as a heavy rainfall that affects the sensors, and, consequently, the driving capability within its operational design domain, the system is designed to hand off control to the driver. If significant danger is detected, the system is designed to come to a safe stop.

NVIDIA follows the V-model (including verification and validation) at every stage of DRIVE development. We also perform detailed analyses of product functionality and related hazards to establish appropriate safety goals. For each identified hazard, we create safety goals to mitigate risk, each rated with an ASIL level. ASIL levels of A, B, C, or D indicate the level of risk mitigation needed, with ASIL D representing the highest level. Meeting these safety goals is the top-level requirement for our design. By applying the safety goals to a functional design description, we create more detailed functional safety requirements.

At the system-development level, we refine the safety design by applying the functional safety requirements to a specific system architecture. Technical analyses—such as Failure Mode and Effects Analysis (FMEA), Fault Tree Analysis (FTA), and Dependent Failure Analysis (DFA)—are applied

iteratively to identify weak points and improve the design. Resulting technical safety requirements are delivered to the hardware and software teams for development at the next level. We've also designed redundant and diverse functionality into our autonomous vehicle system to make it as resilient as possible. This ensures that the vehicle will continue to operate safely when a fault is detected or reconfigure itself to compensate for a fault.

At the hardware-development level, we refine the overall design by applying technical safety requirements to the hardware designs of the board and the SoC. Technical analyses are used to identify any weak points and improve the hardware design. Analysis of the final

hardware design is used to verify that hardware failure-related risks are sufficiently mitigated.

At the software-development level, we consider all software, including firmware. We refine the overall design by applying technical safety requirements to the software architecture. We also perform code inspection, reviews, automated code structural testing, and code functional testing at both unit and integration levels. Software-specific FMEA is used to design better software. In addition, we design test cases for interface, requirements-based, fault-injection, and resource-usage validation methods.

When we have all necessary hardware and software components complete, we integrate and start our

verification and validation processes at the system level. In addition to autonomous vehicle simulation, we also perform system testing and validation.


Safety for Software-Defined Autonomy

Our safety approach is architected for software-defined autonomy. Compared to safety approaches for traditional systems, NVIDIA's safety strategy is:

- > Designed for dynamic system configurations
- > A flexible platform for hardware and software richness
- > Optimized for a growing number of functions
- > Ecosystem-friendly with open system boundaries
- > Designed for AI hardware, software, and tools
- > Expandable with new algorithms
- > Supportive of decomposable safety concepts
- > Designed to handle millions of lines of code
- > Easily updatable over the air
- > Function-aware, data-oriented, and validated
- > Hardware-firmware-software harmonized

All in One: AI Training, Simulation, and Testing

NVIDIA's infrastructure platform includes a data factory to label millions of images. It uses NVIDIA DGX



systems from NVIDIA's own internal cluster for training DNNs, NVIDIA Drive Constellation™ for hardware-in-the-loop simulation, as well as other tools.

AV software development begins with collecting huge amounts of data from vehicles in globally diverse environments and situations. Multiple teams across many geographies access this data for labeling, indexing, archiving, and management before it can be used for AI model training and validation. In addition, real data can be augmented with synthetic data from simulation for scenes that are rare or difficult to label. We call this first step in the autonomous vehicle workflow the “data factory.”

AI model training starts when the labeled data is used to train the models for perception and other self-driving functions. This is an iterative process. The data factory uses the initial models to select the next set of data to be labeled. Deep learning engineers adjust model parameters as needed, and then re-train the DNN, at which point the next set of labeled data is added to the training set. This process continues until the desired model performance and accuracy are achieved.

Self-driving technology must be evaluated again and again during development in a vast array of driving conditions to ensure that the vehicles are far safer than human-driven vehicles. Simulation runs test-drive scenarios in a virtual world, providing rendered

sensor data to the driving stack and carrying out driving commands from the driving stack. Re-simulation plays back previously recorded sensor data from the real world to the driving stack. The AI model is then validated against a large and growing collection of test data.

Hardware

The NVIDIA DRIVE AGX hardware architecture is scalable, covering everything from entry-level advanced driver assistance systems to fully autonomous robotaxis. The current-generation DRIVE AGX Orin SoC's safety architecture was developed over several years by hundreds of architects, designers, and safety experts based on analysis of hundreds of safety-related modules.

Built as a software-defined platform, NVIDIA DRIVE AGX Orin is developed to enable architecturally compatible platforms that scale from a Level 2 to a full self-driving Level 5 vehicle, enabling OEMs to develop large-scale and complex families of software products. All of NVIDIA's DRIVE AGX SoC product family (DRIVE AGX Thor, Orin, and Xavier™) are programmable through open CUDA and TensorRT APIs and libraries, so developers can leverage their investments across multiple product generations.

In 2022, NVIDIA introduced DRIVE AGX Thor, our next-generation centralized computer for safe and secure autonomous vehicles. It features 8-bit floating point support (FP8), delivering an unprecedented 1,000 INT8 TOPS/1,000 FP8 TFLOPS/500 FP16 TFLOPs of performance.

This next-gen AV processor unifies intelligent functions—including advanced driver assistance and in-vehicle infotainment—into a single architecture for greater efficiency, safety, and security.

It also comes packed with cutting-edge AI capabilities and will integrate the new NVIDIA Blackwell GPU architecture, designed for transformer, LLM, and generative AI workloads.

Available for automakers' 2025 models, DRIVE AGX Thor will accelerate production roadmaps by bringing higher performance and advanced features to market at the same time.

Software

The Perception module in the DRIVE SDK takes sensor data and uses a combination of deep learning and traditional signal processing to determine an understanding of the vehicle's environment—called the World Model. Once the environment is understood,

the Planning module uses this information to find and score a set of trajectories and determine the best route. The Vehicle Dynamics Control module can then transform the chosen path into vehicle actuation.


DRIVE SDK currently uses more than 20 DNN models running simultaneously, in addition to a large suite of computer vision and robotics algorithms. However, the number of DNNs and the capabilities they cover are continually growing.

Each major function—such as sensor processing, AI-based perception, localization, trajectory planning, and mapping—is performed with multiple redundant and diverse methods to achieve the highest level of safety. For example, the DRIVE SDK uses embedded modules for detecting and handling obstacles and drivable space. For wait conditions, we detect traffic lights, stop signs, intersections, and stop lines. For paths, we detect lane edges and drivable paths. This detection is happening over multiple frames, and objects are tracked over time.

We also layer diversity by using multiple sensor types (radar, camera, lidar, and ultrasonic). The combination of diverse DNNs, tracking of objects over multiple frames, and presence of different sensor types ensures safe operation within the operational design domain. Additionally, the integrated functional safety mechanisms enable safe operation in the event of a system fault.

Vehicles and Sensors


DRIVE AGX Hyperion is a reference vehicle implementation of the DRIVE AGX platform to enable self-driving development, data collection and ingestion, verification, and validation across automation levels. The platform leverages multiple sensor modalities—including cameras, radars, lidars, IMUs, and ultrasonic sensors—and is deployable to a variety of vehicle types.



Data Center

After collecting sensor data, we process it and, in the case of camera data, select images to be labeled for training the AI. The entire process is continuously validated. We label not only objects and images within captured frames, but also scenarios and conditions in video sequences. The more diverse and unbiased data we have, the safer the DNNs become. We also define key performance metrics to measure the collected data quality and add synthetic data into our training datasets using NVIDIA Omniverse Replicator. This lets developers generate pre-labeled ground-truth data to bootstrap algorithm development. The ultimate goal is to continuously add training data to build a comprehensive matrix of locations, conditions, and scenarios. Performance of neural network models is validated using independent test data and retested as models are trained on new data.

In addition to labeling the objects in an image, we label the conditions under which data was collected. This provides a matrix of conditions we can use as a test dataset to test the performance



of our DNN models against a wide range of scenarios, weather conditions, and times of day.

GPUs in the data center are used extensively to investigate new DNNs with diverse datasets,

continually train neural network models, analyze the results of workflows, and test and validate outcomes using large-scale systems for simulation in virtual worlds and replay of collected data.

On-Road Testing

NVIDIA created the DRIVE Road Test Operating Handbook to ensure a safe, standardized on-road testing process. This document specifies what must be done before, during, and upon completion of every road test. As recommended in the U.S. DOT report *Preparing for the Future of Transportation: Automated Vehicles 4.0*¹², NVIDIA's process is modeled on the FAA-certified *Pilot's Operating Handbook* that must be carried in-flight with every general aviation aircraft in the U.S.

On-road testing is always performed with a highly trained safety driver continuously monitoring the vehicle's behavior and ready to immediately intervene when necessary. A test operator may also be in the vehicle and monitors the self-driving software—for instance, checking that the objects detected by the car correspond to those viewed live—and that the vehicle's path is valid for current road conditions.

Before allowing software to be tested on the road, it's extensively tested using unit tests, integration tests, and system simulation.



Developer Training and Education



NVIDIA is committed to making developer education easily accessible, helping both experts and students learn more about these breakthrough technologies. The NVIDIA Deep Learning Institute (DLI) offers multiple courses on how to design, train, and deploy DNNs for autonomous vehicles. We also produce a wide range of content to answer common questions and now have over two million registered developers across eight different domains, such as deep learning, accelerated computing, autonomous machines, and self-driving cars.

In addition, NVIDIA hosts the GTC conference to help educate students, developers, and executives on accelerated computing, AI, and autonomous vehicles. Each conference features hundreds of sessions, panels, and hands-on courses, as well as groundbreaking technology demos and partner exhibits. Each conference begins with the Keynote from CEO Jensen Huang and features hundreds of sessions, panels, and hands-on courses, as well as groundbreaking technology demos and partner exhibits.

Summary

NVIDIA is unique in providing foundational technologies for the design, development, and manufacture of safe and reliable software-defined autonomous vehicles. Our ability to combine the power of visual and high-performance computing with artificial intelligence and proven software development makes us an invaluable partner to vehicle manufacturers and transportation companies around the world.

We adhere to the industry's most rigorous safety standards in the design and implementation of the powerful NVIDIA DRIVE AGX platform, and we collaborate with industry experts to address current and future safety issues. Our platform aligns with and supports the safety goals of the major autonomous vehicle manufacturers and robotaxi companies.

Building safe autonomous vehicle technology is one of the largest, most complex endeavors our company has ever undertaken. We've invested billions of dollars in research and development, and many thousands of engineers throughout the company are dedicated to this goal. As of now, more than 1,500 engineer-years have been invested in our automotive safety process.

There are currently more than 80 AV companies with test vehicles on the road powered by NVIDIA technology. They recognize that greater compute in the vehicle enables redundant and diverse software algorithms to deliver increased safety for every driver.

We fundamentally believe that self-driving vehicles will bring transformative benefits to society. By eventually removing human error from the driving equation, we can prevent the vast majority of accidents and minimize the impact of those that do occur. We can

also increase roadway efficiencies and curtail vehicle emissions. Finally, people who may not have the ability to drive a car will gain the freedom of mobility when they can easily summon a self-driving vehicle.

NVIDIA holds a key role in the development of AVs that will revolutionize the transportation industry over the next several decades. Nothing is more exciting to us than overcoming technology challenges, making people's lives better and our roads safer.



Appendices

Appendix A: NVIDIA Activity in Expert Groups

NVIDIA is respected as an organization of experts in relevant fields as demonstrated by the active and leadership role that our experts have in international standardization working groups. Working groups benefiting from our expertise include:

- > ISO TC 22/SC 32/WG 8, ISO 26262, “Functional Safety” and ISO 21448 “Safety of the intended functionality”
- > ISO TC 22/SC 32/WG 13, ISO TS 5083, “Safety and cybersecurity for automated driving systems—Design, verification and validation methods”
- > ISO TC 22/SC 32/WG 14, ISO PAS 8800, “Safety and Artificial Intelligence”
- > ISO TC 22/SC 32 and SAE TEVEES 18A, ISO/SAE 21434, “Cybersecurity engineering”
- > ISO/TR 9839, “Application of Predictive Maintenance with ISO 26262-5”
- > IEC 61508, “Functional safety of electrical/electronic/programmable electronic safety-related systems”
- > IEEE 2846-2022, “A Formal Model for Safety Considerations in Automated Vehicle Decision Making”
- > IEEE P2851, “Standard for functional safety data format for interoperability within the dependability lifecycle”
- > IEEE Computer Society Functional Safety Standards Committee (FSSC)

- > ISO/IEC JTC1 SC42 JWG4, ISO/IEC TR 5469 and ISO/IEC TS 22440, “Artificial intelligence — Functional safety and AI systems”
- > Euro NCAP, through the European Association of Automotive Suppliers
- > Comité de Liaison de la Construction d’Equipements et de Pièces d’Automobiles
- > UNECE Working Groups on Functional Requirements for Automated Vehicles (FRAV) and Validation Methods for Automated Driving (VMAD)
- > UNECE Working Groups on Dynamic Control Assistance System (DCAS)
- > SAE Committees on Automotive Functional Safety and Automotive Ground Vehicle AI
- > Multiple worldwide R&D consortia, technical review committees, and R&D chair roles

Appendix B: National and International Safety Regulations and Recommendations

NVIDIA adheres to national and international safety recommendations, including:

International Organization for Standardization (ISO)

We adhere to ISO 26262 and ISO 21448 (SOTIF). ISO 26262:10 addresses functional safety in road vehicles. We apply ISO 26262 to the application, middleware, operating system, board, and chip levels. ISO 21448:11 addresses safety of the intended functionality in road vehicles. It reuses and extends the ISO 26262 development process to address SOTIF concerns. We also follow the ongoing standardization efforts on AI safety in ISO PAS 8800, ISO/IEC TR 5469, and ISO/IEC TS 22440.

Global NCAP (New Car Assessment Program)

Regional NCAPs adjust safety practices to their particular markets, and NVIDIA will evaluate performance with all local NCAPs. The European New Vehicle Assessment Program (Euro NCAP) provides consumers with an independent safety assessment of vehicles sold in Europe. Euro NCAP published its 2025 Roadmap¹², which presents a vision and strategy to emphasize primary, secondary, and tertiary vehicle safety. We’re currently addressing these Euro NCAP recommendations:

- > Automatic Emergency Steering
- > Pedestrian and Cyclist Safety
- > Assisted Driving Testing
- > Simulation and Assessment
- > Child Presence Detection
- > Autonomous Emergency Braking
- > Cybersecurity
- > V2X
- > Driver Monitoring
- > Human-Machine Interface (HMI)
- > Pedestrian and Cyclist Safety
- > Simulation and Assessment
- > Child Presence Detection
- > Cybersecurity

Appendices

Appendix C: NHTSA Elements of Safety

The National Highway Traffic Safety Administration outlined key topics for self-driving safety in its [Automated Driving Systems 2.0: A Vision for Safety](#). Of the 12 safety elements representing industry consensus on safety for the use of automated driving systems on public roadways, 10 are relevant to NVIDIA.

System Safety

NVIDIA has created a system safety program that integrates robust design and validation processes based on a systems-engineering approach with the goal of designing automated driving systems with the highest level of safety and free of unreasonable safety risks.

Object and Event Detection and Response

Object and event detection and response refers to the detection of any circumstance that's relevant to the immediate driving task, and the appropriate driver or system response to this circumstance. The NVIDIA DRIVE AV module is responsible for detecting and responding to environmental stimuli, both on and off the road. The NVIDIA DRIVE IX module helps monitor the driver and take mitigation actions when they're required.

Operational Design Domain

NVIDIA has developed an extensive set of operational design domains for individual driving automation systems or features, as recommended by NHTSA. Each operational design domain includes the following information at a minimum to define the product's capability boundaries: roadway types, geographic area and geo-region conditions, speed range, environmental conditions (weather, time of day, etc.), and other constraints.

Fallback (Minimal Risk Condition)

Our products enable the vehicle to detect a system malfunction or breach of the operational design domain, and then transition the system to a safe or degraded mode of operation based on a warning and degradation strategy. Every NVIDIA autonomous driving system includes a fallback strategy that enables the driver to regain proper control of the vehicle or allows the autonomous vehicle to return to a minimal risk condition independently. Our HMI products can be used to notify the driver of a potentially dangerous event and return the vehicle to a minimal risk condition independently, or alert the

driver to regain proper control. The minimal risk conditions vary according to the type and extent of a given failure.

Validation Methods

Validation methods establish confidence that the autonomous system can accomplish its expected functionalities. Our development process contains rigorous methods to verify and validate our products' behavioral functionality and deployment. To demonstrate the expected performance of an autonomous vehicle for deployment on public roads, our test approaches include a combination of simulation, test track, and on-road testing. These methods expose performance under widely variable conditions, such as when deploying fallback strategies.

Human-Machine Interface

DRIVE IX provides an open software stack for cockpit solution providers to build and deploy features that will turn personal vehicles into interactive environments, enabling intelligent assistants, graphic user interfaces and immersive media and entertainment.

Vehicle Cybersecurity

NVIDIA employs a rigorous security development lifecycle into our system design and hazard analysis processes, including threat models that cover the entire autonomous driving system—hardware, software, manufacturing, and IT infrastructure. The NVIDIA DRIVE AGX platform has multiple layers of defense that provide resiliency against a sustained attack.

NVIDIA also maintains a dedicated Product Security Incident Response Team that manages, investigates, and coordinates security vulnerability information internally and with our partners. This allows us to contain and remediate any immediate threats while openly working with our partners to recover from security incidents.

Data Recording

NVIDIA replay enables real data from sensors placed on test vehicles that are driving on public roads to be fed into the simulation. To maximize the safety of self-driving vehicles, NVIDIA offers a combination of simulated data to test dangerous road scenarios coupled with real-world data from replay.

Consumer Education and Training

We continually develop, document, and maintain material to educate our employees, suppliers, customers, and end consumers. We offer multiple AI courses under our [Deep Learning Institute \(DLI\)](#) and we report on new knowledge and developments at the [NVIDIA GTC](#) around the world. We also collaborate with research organizations to invent improved approaches to autonomy and maintain the highest integrity level to co-create world-class thought leadership in the autonomous vehicle domain.

Federal, State, and Local Laws

We operate under the principle that safety is the first priority and comply with international, federal, state, and local regulations, as well as safety and functional safety standards. We also frequently communicate with regulators to ensure that our technology exceeds all safety standards and expectations. We're active in standardization organizations to advance the future of autonomous driving.

Appendix D: References

1. NHTSA Automated Driving Systems: <https://www.nhtsa.gov/vehicle-manufacturers/automated-driving-systems>
2. ISO 26262: Road Vehicles—Functional Safety. The International Organization for Standardization (ISO), 2018. Part 1: Vocabulary: Terms and Definitions: <https://www.iso.org/standard/68383.html>
3. ISO 21448, Road vehicles—Safety of the intended functionality, 2022: <https://www.iso.org/standard/77490.html>
4. ISO/CD PAS 8800: ISO/CD PAS 8800—Road Vehicles; Safety and artificial intelligence
5. ISO/IEC TR 5469:2024: ISO/IEC TR 5469:2024—Artificial intelligence; Functional safety and AI systems
6. ISO/IEC AWI TS 22440-1/-2/-3: ISO/IEC AWI TS 22440-1—Artificial intelligence; Functional safety and AI systems, Part 1: Requirements
7. IEEE 2846-2022 web page: [IEEE SA - IEEE 2846-2022](https://www.ieee.org/standards/2846-2022)
8. IEEE P2851 web page: <https://sagroups.ieee.org/2851/>
9. NIST Cybersecurity: <https://www.nist.gov/topics/cybersecurity>
10. GDPR: <https://gdpr.eu/>
11. UNECE Regulation 155: <https://unece.org/transport/documents/2021/03/standards/un-regulation-no-155-cyber-security-and-cyber-security>
12. Preparing for the Future of Transportation: Automated Vehicles 4.0 (AV 4.0), 2020: <https://www.transportation.gov/av/4>

Notice

All information provided in this report, including commentary, opinion, NVIDIA design specifications, drawings, lists, and other documents (together and separately, "materials") are being provided "as is." NVIDIA makes no warranties, expressed, implied, statutory, or otherwise with respect to materials, and expressly disclaims all implied warranties of noninfringement, merchantability, and fitness for a particular purpose.

© 2026 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, Cosmos, CUDA, DGX, DRIVE, DRIVE AGX, DRIVE AGX Hyperion, DRIVE AGX Orin, DRIVE AGX Thor, DRIVE AGX Xavier, DRIVE Constellation, Omniverse, Omniverse Cloud Sensor RTX, OVX, RTX, and TensorRT are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. 4685112. MAR26

