



## Introduction

# Table of contents

## AV 2.0: Making Vehicles Safer With AI

### The Four Pillars of Safe Autonomous Driving

Pillar 1: Artificial Intelligence Design and Implementation Platform

Pillar 2: Development Infrastructure for Deep Learning

Pillar 3: Physically Accurate Sensor Simulation for AV Development

Pillar 4: Best-In-Class Pervasive Safety and Cybersecurity Program

### Architected for Safety

Hardware

Software

Vehicles and Sensors

Data Center

On-Road Testing

Developer Training and Education

### Summary

# AV 2.0: Making Vehicles Safer With AI

## AV 2.0: Making Vehicles Safer With AI

Building an autonomous system to navigate safely in the complex physical world is challenging. The system needs to perceive and understand its surrounding environment holistically, and then make correct, safe decisions in a fraction of a second. This requires human-like situational awareness to handle potentially dangerous or rare scenarios.

### **AV 2.0 and End-to-End Driving**

AV software development has traditionally been based on a modular approach, with separate components for object detection and tracking, trajectory prediction, and path planning and control. A new era of autonomous vehicle technology, known as AV 2.0, has emerged. AV 2.0 is characterized by large, unified AI models that can control multiple parts of the vehicle stack, from perception and planning to control. In contrast to AV 1.0's focus on refining a vehicle's perception capabilities using multiple deep neural networks, AV 2.0 calls for comprehensive in-vehicle intelligence to drive decision-making in dynamic, real-world environments using an approach known as "end-to-end driving." End-to-end autonomous driving systems use a unified model to take in sensor input and produce vehicle trajectories. This helps avoid overcomplicated pipelines and provide a more holistic, data-driven approach to handle real-world scenarios.

### **A Focus on Safety**

AV 2.0 will play an important role in building and validating safe AV systems. Examples of how NVIDIA technology is being applied in this area include:

**1. Simulation:** Safe AV systems must be prepared to navigate rare and unusual situations safely. NVIDIA is developing capabilities to create high-quality, realistic simulations of both traffic and sensors, and construct counterfactual situations guided by natural language descriptions of safety-critical scenarios. During development, these capabilities will allow augmenting real-world training data to improve the robustness of AV modules. At evaluation time, they can also offer an additional mechanism to validate AV systems at scale, complementing realworld testing and validation.

**2. Safe interaction:** As AV systems are deployed on our roads, they must interact with human road users. NVIDIA is taking advantage of AI to learn predictive models of driving behavior, and to use these predictions to understand the impact of an AV's actions on other road users. Using these capabilities, developers can design AV systems that interact responsibly with other drivers and pedestrians and minimize the risk of accidents.

**3. Anomaly detection:** AV safety requires systems that can handle anomalous situations

reliably. AI models that predict scene evolution can enable systems to evaluate which anomalies might have safety-critical impacts and require execution of fail-safe behavior versus which anomalies can be safely ignored. NVIDIA is exploring how learned future prediction models can be used to evaluate the risk of perception failures.

## **The Ultimate Triathlon**

The race to develop safe self-driving cars isn't a sprint, but more a never-ending triathlon, with three distinct yet crucial parts operating simultaneously: AI training, simulation, and autonomous driving. Each requires its own accelerated computing platform. Together, the purpose-built, full-stack systems form a powerful triad that enables continuous development cycles, always improving in performance and safety. A model is first trained on an AI supercomputer such as NVIDIA DGX™. It's then tested and validated in simulation—using the NVIDIA Omniverse™ platform and running on an NVIDIA OVX™ system—before entering the vehicle, where, lastly, the NVIDIA DRIVE AGX™ platform processes sensor data through the model in real time using [NVIDIA DriveOS™](#), the operating system for safe, AI-defined autonomous vehicles. AV 2.0 shows great promise in building and validating safer AV systems. As with any AI system, it's important to use it responsibly. We advocate augmenting generative AI systems with high-quality uncertainty quantification and guardrails. Using these capabilities, AVs can navigate the complex and unpredictable world more safely and reliably. Watch a video about the [Hydra-MDP model](#) from NVIDIA Research, winner of the CVPR 2024 Autonomous Grand Challenge for End-to-End Driving

## **The Four 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.

- > Pillar 1: Artificial Intelligence Design and Implementation Platform
- > Pillar 2: Development Infrastructure to Support Deep Learning
- > Pillar 3: Physically Accurate Sensor Simulation for AV Development
- > Pillar 4: Best-In-Class Pervasive Safety and Cybersecurity Program

### **Pillar 1: 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 on-board supercomputers capable of processing all the sensor data in real time.

## **NVIDIA DRIVE Hardware**

Our underlying hardware solutions include:

NVIDIA DRIVE Hyperion (3840x2160, AR: 1.78) **NVIDIA DRIVE AGX Hyperion**

NVIDIA DRIVE AGX Hyperion™ is an end-to-end, modular reference architecture for designing autonomous vehicles (AVs). 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 Orin (600x340, AR: 1.76) **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 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 Thor (3840x2160, AR: 1.78)**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 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 (3840x2160, AR: 1.78)**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 (2394x1344, AR: 1.78)**NVIDIA DRIVE Software Development Kits**

Software is what turns a vehicle into an intelligent machine. The open NVIDIA DRIVE SDK gives developers all the building blocks and algorithmic stacks needed for autonomous driving. It empowers developers to efficiently 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 DriveOS, the first safe operating system for in-vehicle accelerated computing. 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 DriveOS that are fundamental to autonomous vehicle development. These consist of the sensor abstraction layer (SAL) and sensor plugins, data recorder, vehicle I/O support, and a DNN framework. It's modular, open, and designed to be compliant with automotive industry software standards.
- NVIDIA offers an AI-assisted driving platform that can handle both highway and urban traffic with the utmost safety. It can use the high-performance compute and sensor set of NVIDIA DRIVE AGX 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.

## **Pillar 2: 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.



## **Pillar 3: 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 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, built on OpenUSD and NVIDIA RTX™, are 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 Cloud Sensor RTX microservices enable physically based and neural rendering of sensors commonly deployed on autonomous vehicles, including cameras, lidar, radar, and ultrasonics sensors. The rendered synthetic data and ground-truth labels can be used for training perception models and validating the AV software stack in closed-loop testing.

The neural reconstruction engine is a new AI toolset for the AV simulation platform that uses multiple AI networks to turn recorded sensor data into usable world models for simulation. The new pipeline uses AI to automatically extract the key components needed for simulation, including the environment, 3D assets, and scenarios. These pieces are then reconstructed into simulation scenes that have the realism of data recordings, but are fully reactive and can be manipulated as needed. Achieving this level of detail and diversity by hand is costly, time consuming, and not scalable.

In addition, fVDB is a new open-source deep-learning framework that can be used to generate large-scale scenes for training AVs using real-world 3D data. It builds AI operators on top of OpenVDB to create high-fidelity virtual representations of real-world environments. These rich 3D datasets are AI-ready for efficient model training and inference. Soon, fVDB functionality will be available as NVIDIA NIM microservices that enable developers to incorporate the fVDB core framework into Universal Scene Description (OpenUSD) workflows. fVDB NIM microservices generate OpenUSD-based geometry in NVIDIA Omniverse.

## **Pillar 4: 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.<sup>1</sup>

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)<sup>2</sup> must also be designed and tested to be performant on all safety-critical metrics related to the intended functionality (ISO 21448)<sup>3</sup>.

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 that 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)<sup>7</sup> and IEEE P2851 (on Exchange/Interoperability Format for Safety Analysis and Safety Verification of IP, SoC, and Mixed Signal ICs<sup>8</sup>).

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.

NVIDIA uses TÜV SÜD—an independent, accredited assessor—to ensure compliance with the International Organization for Standardization (ISO) 26262:2018 Functional Safety Standard for Road Vehicles.

NVIDIA DRIVE platforms and processes have recently been certified and assessed by TÜV SÜD:

- NVIDIA DRIVE core development processes are certified as ISO 26262 Automotive Safety Integrity Level (ASIL) D compliant.
- The NVIDIA DRIVE 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 Orin-based system, which unifies

the DRIVE Orin SoC and DRIVE AGX Orin board,

completed concept assessment and is deemed to

meet ISO 26262 ASIL D requirements.

- Development of NVIDIA DriveOS 6.x is in progress

and will be assessed by TÜV SÜD for ASIL D compliance. This follows the certification of change to DriveOS 5.2 ASIL B certification, which includes NVIDIA CUDA libraries and the NVIDIA TensorRT software development kit for real-time

AI inferencing.

## Cybersecurity

An AV 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 established 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 adhere to standards set by NIST (National Institute of Standards and Technology)<sup>9</sup> and GDPR (General Data Protection Regulations)<sup>10</sup> 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. 15511. In addition, we use the ISO/SAE 21434 cybersecurity process and align our automotive development practices accordingly to make compliance claims easier besides leveraging processes and practices from other cybersecurity-sensitive industries with ISA/IEC 62443.

We participate in the SAE J3101 standard development, which ensures 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, including end-to-end traceability on security requirements, threat models that cover the entire autonomous driving system. This includes hardware, software, manufacturing, and IT infrastructure, ensuring secure design and coding guidelines are in place. The DRIVE platform also 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. This includes delving 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.

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, therefore, 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 our products' functionality and related hazards to develop safety goals for the product. 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, 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, Self-Driving Safety Report | 18 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 testdrive scenarios in a virtual world, providing rendered sensor data to the driving stack and carrying out driving commands from the driving stack. Resimulation 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 Orin SoC’s safety architecture was developed over several years by hundreds of architects, designers, and safety experts based on analysis of hundreds of safetyrelated 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 SoC product family (DRIVE 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 Thor, our nextgeneration 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 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.

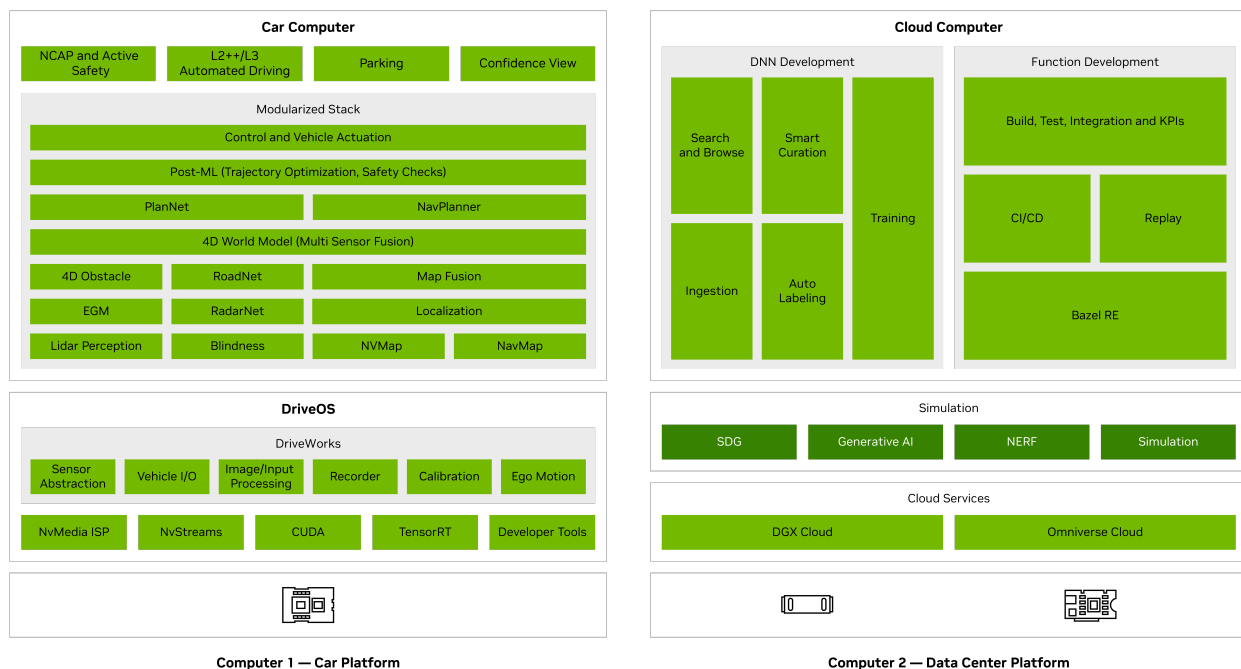
We also layer diversity by using multiple sensor types (radar, camera, lidar, and ultrasonic).



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 Hyperion is a reference vehicle implementation of the DRIVE 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 4.012*, 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 is also 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.

We can also modify our processes when it's not possible to test in-person. The NVIDIA teleoperation system enables the human co-pilot to remotely monitor the vehicle, while the virtual testing platform makes it possible to safely and securely test vehicles virtually.

Before allowing software to be tested on-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 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.





**Notice**

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, 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 THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A

PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.<br/><br/><br/><br/><b>Trademarks</b><br/><br/><br/>NVIDIA and the NVIDIA logo 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.<br/>

© Copyright 2024, NVIDIA. PDF Generated on 03/05/2025