NVIDIA Response to NHTSA ANPRM DOT Docket 2020-0106
24 March 2021

Introduction

NVIDIA is pleased to provide comments in response to DOT Docket Number: NHTSA-2020-0106.

The development of Automated Driving Systems (ADS) is one of the great technological challenges of our time. Successful deployment of ADS will deliver great benefit to society, but its full potential will not be realized without partnership among researchers, developers, manufacturers, and regulators.

NVIDIA envisions a far future of SAE J3016 Level 5 (unrestricted Operational Design Domains (ODD)) and high-capability Level 4 (minimally-restricted ODD) million- to billion- ADS vehicle fleets, enabled by mobile Artificial Intelligence (AI) supercomputing, to serve a wide variety of transportation system needs, delivered by many diverse manufacturers. The average ADS will have safety performance that is substantially better than the statistically average human driver.

Nothing in this world is perfect, and as such perfect ADS safety may never be realized, even in the long run. But the hundred-year history of passenger airline travel shows that extremely high levels of safety can be achieved when every incident is regarded as an opportunity to learn. Cooperation among stakeholders, guided by well-designed regulation frameworks, will be vital to achieving and preserving progress and relentlessly advancing safety for all.

Our shared task is to envision a framework that will evolve together with industrywide ADS advancement from today's state of the art to the far future of driving automation safety maturity.

Overall ADS dependability rests on three foundational pillars – functional safety, safety of the intended functionality (SOTIF), and resilience against attack (cybersecurity). Functional safety seeks to minimize harm due to malfunctioning hardware, software, or systems. SOTIF is a complementary aspect of safety – it seeks to minimize the risk of harm due to limitations in sensing and algorithms when there is no malfunction. Functional safety is a relatively mature discipline, with decades of experience in industrial and automotive settings. Standards for ADS SOTIF are only now emerging; SOTIF remains an area of active research in industry and academia.

The expectations for ADS safety will evolve over time. Initially it will be enough to avoid ego-vehicle at-fault collisions and react appropriately when other actors' behaviors cause hazardous events. In the long run, ADS will exhibit higher levels of situational awareness, taking account of the mental state of other actors - their emotions, knowledge, and intentions. For example, an ADS vehicle being passed aggressively on a two-lane road with a rise ahead will be aware that the passer does not have enough visibility, and will anticipate, not simply react to, the other vehicle swerving in panic when the oncoming car suddenly appears. The more safe that ADS fleets become, the more important the handling of the rare situations will be to their continued progress. ADS vehicles will not be allowed to enter service unless they meet a very high level of safety. With ongoing research, innovation and accumulated experience, ADS vehicles will safely handle increasingly rare and complex situations, further improving safety for themselves and for those around them.

While allowing for flexibility in implementation, we propose that ADS should have certain fundamental attributes in common. When the state of the art has reached sufficient maturity and consensus, agreed-upon fundamentals may serve as a basis for rulemaking.

1. **ADS shall have on-board monitoring to confirm safe behavioral planning.** Primary behavioral planning will need to cope with a vast state space of possible conditions. Over the long run this may come to include semantic, cognitive, situational awareness, far beyond today's elementary driving behaviors. It will be intractable, if not impossible, to prove with absolute certainty ahead of time that these sophisticated planners will behave safely in every possible situation they will encounter. Therefore, safe ADS must include a monitoring subsystem whose purpose is to prevent unsafe control actions initiated by the primary planner from being carried out. Such a monitoring subsystem will be designed robustly and simply enough so that its independent collision avoidance capability can be verified. Any incident in a development or production fleet of the monitor intervening will constitute a reportable event. NVIDIA's Safety Force Field (https://www.nvidia.com/content/dam/en-zz/Solutions/self-driving-cars/safety-force-field/an-introduction-to-the-safety-force-field-updated.pdf) can serve as this "virtual guardrail" and is openly available to the ADS industry.

2. **Safe autonomous operation of ADS shall only rely on local resources.** ADS in control of moving vehicles shall not rely on real-time data communications external to the vehicle, which are subject to interruption, to safely perform their control task. Restated, only on-board data should be necessary for safe operation. This principle does not preclude receiving signals from outside the vehicle, whether audio, visual, or by radio in order to comply with the directions of first responders. Nor does it preclude the non-real time download of control or map updates or the upload of incident reporting. Vehicle-to-vehicle or vehicle-to-infrastructure ("V2X") may be used if available, but availability of V2X should not be relied upon to achieve minimum acceptable safety. Any remote control of a vehicle must be by an authorized operator and be robust against communication disruption and against cyberattack.

3. **ADS shall use combinations of diverse test methods to demonstrate safety performance.** ADS development, validation, and ongoing assessment will use a combination of high-fidelity simulation, controlled test track testing, limited public-road testing, and deployed fleets of vehicles. A system being evaluated for release should be expected to pass all test cases during simulation or controlled test-track testing; any allowed waivers should be justified by careful risk analysis. Zero real-world failures will be difficult to guarantee at deployment, but each supplier must strive to demonstrate very low frequency of at-fault failures before deployment, and then to learn and improve from every real-world failure. Tolerable field failure rates will diminish over years and decades as the industry develops and learnings accumulate.

4. **Physical testing of a complete vehicle as a "black box" is insufficient to demonstrate ADS safety.** Physical vehicle track and road testing, while necessary, is by no means sufficient to

empirically demonstrate the safety performance of an ADS. Physical testing can practically demonstrate hundreds or thousands of situations, at best. Simulations, in contrast, can demonstrate millions to billions of scenarios. Yet even billions of merely randomly chosen scenarios cannot find everything -- we need ongoing advancement of a large corpus of directed stress-case simulation test scenarios built from ongoing road experience. An efficient automated reporting path from fleet operation of collision, incident, unexpected, or otherwise anomalous events will always be needed, to continuously augment a standard qualification simulation suite. Due to the complexity of advanced ADS, qualitative "white box" examination of element design, performance, and AI training methodology by expert assessors will be an important cross-check against the objective demonstration measures.

5. **Rapid incident reporting and response infrastructure will be needed.** The timely discovery of safety vulnerabilities will depend on the ability to identify patterns or clusters of adverse events, on timescales ranging from hours to days or weeks after release of a new ADS major or incremental update. Rapid recovery from incidents or discovered safety vulnerabilities will require robust low-latency methods to update ADS control systems in the field, balancing the need for rapidly "pushing" an improved control system or a temporary ODD limitation against the requirement to thoroughly confirm that the update is in fact a safety improvement and not a regression.

Question 1. Describe your conception of a Federal safety framework for ADS that encompasses the process and engineering measures described in this notice and explain your rationale for its design.

The framework should leverage industry standards that have achieved broad consensus, and have clearly defined scope within the domains of the functional safety or SOTIF disciplines, e.g. ISO 26262, ISO 21448, and cybersecurity, ISO/SAE 21434. Additional standards which provide guidance include UL 4600, ISO/TR 4804 and others in varying stages of proposal or drafting.

NHTSA can model compliance and assessment activities after the FAA DER (Designated Engineering Representative) "super-assessor" model of deep expert independent assessment. This is particularly appropriate for safety case and white box qualitative reviews. Suppliers should be permitted to perform internal independent self-assessment as long as appropriate safety processes are followed, a practice that ISO 26262 has described.

The safety framework should account for non-traditional automotive domains where they apply to the safety of software-defined ADS behaviors, for example integrity of Over-the-Air (OTA) software updates, incident reporting, and cybersecurity e.g., ISO/AWI 24089 or ISO/SAE 21434. NHTSA may need to obtain expanded regulatory authority for these new domains.

The framework should evolve towards a steady-state model which incorporates learnings from production ADS fleets into regular updates of simulation-test, track-test and road-test validation and assessment suites. The target safety metrics should be derived from baselines established by the aggregate safety statistics of human drivers, classified by degree of difficulty of the driving situations.

Question 2. In consideration of optimum use of NHTSA's resources, on which aspects of a manufacturer's comprehensive demonstration of the safety of its ADS should the Agency place a priority and focus its monitoring and safety oversight efforts and why?

NHTSA can fulfill an important role by collecting existing incident, collision, and fatality data and classifying such data by, for example, traffic conditions (ranging from open rural highway to congested city rush-hour), environmental conditions (weather, lighting), and road conditions (construction activity, road debris) to establish situation-appropriate safety performance benchmarks.

Demonstration of safety occurs during different phases in development of a given ADS – functional safety and SOTIF, from concept through research, prototype, hardware-in-the-loop simulation validation, test-track, public road with initially-limited and later expanded ODDs. In early phases, safety demonstrations rely more heavily on argumentation (e.g., safety case) and extrapolation; in middle and mature phases, actual fleet operation statistics have an increasing role, viewed at multiple time scales, e.g., daily, monthly, yearly. A common reporting framework will be needed to compare different manufacturers' results against requirements. Requirements themselves will evolve over time.

If NHTSA chooses to leverage the Designated Engineering Representative model, the Agency should develop comprehensive training to establish authority, scope, and practices to guide the work of these independent experts.

It will be challenging to monitor fleet-wide flows of nominal and off-nominal events without active online automation and filtering. NHTSA can encourage industry associations to cooperatively develop standards which specify the characteristics of event transmission. These collaborations can arrive at consensus on what type of filtering/interpretation is appropriate to occur in real-time on each vehicle, versus in near-real-time on datacenters managed by the manufacturers and/or datacenters managed by NHTSA.

Question 3. How would your conception of such a framework ensure that manufacturers assess and assure each core element of safety effectively?

Quantitative safety metric targets, properly organized and classified by Operational Design Domain, will enable manufacturers to track and report real-world fleet performance in a way that can be compared against their own performance over time, against that of other manufacturers, and against absolute objective performance as established and updated from time to time by NHTSA.

Before fleet statistics are available, i.e., pre-deployment, manufacturers should make available summarized qualitative safety case evidence and argumentation, including development, test flows, and qualification activities, following the functional safety lifecycle, in a form reviewable by NHTSA or its authorized DERs. Objective evidence can include simulated test results on curated suites of stress-case scenarios.

Question 4. How would your framework assist NHTSA in engaging with ADS development in a manner that helps address safety, but without unnecessarily hampering innovation?

The framework should allow each developer the flexibility to define and justify their own development against a core set of metrics, using the safety case approach of evidence and argumentation. While qualitative reviews can give insight into development processes, NHTSA must take care not to delay or withhold approval based solely on subjective opinions. Past legal precedents have established that NHTSA shall evaluate safety using objective metrics. Practitioner training e.g., "handbooks" for Independent Assessors who act under NHTSA authority must focus on the need for objectivity.

To balance safety and innovation, restricted interim approvals can limit short-term public safety risks in exchange for the longer-term public safety benefits of accelerating ADS maturity.

Consensus standards not yet in existence may need to be created. Industry and professional associations should be enlisted to identify and develop appropriate standards, in response to calls by NHTSA.

Question 5. How could the Agency best assess whether each manufacturer had adequately demonstrated the extent of its ADS' ability to meet each prioritized element of safety?

Manufacturers should make their safety case argumentation available for review by the Agency or its DERs. This includes verification of the safety monitoring subsystems within the ADS, high-fidelity simulation and test-track empirical results and methodologies, and public road test and deployment results, including actual incident occurrences and anomalies as recorded by embedded safety monitoring subsystems. Raw data containing these results will be enormous and would be impractical to transmit in its entirety between manufacturers and NHTSA. The Agency can help by defining, or encouraging industry to define by consensus, standards for the capture, retention, filtering, and anonymization of such data. It will be important that the assessment processes safeguard manufacturers' proprietary data to prevent unauthorized release.

Question 6. Do you agree or disagree with the core elements (i.e., "sensing," "perception," "planning" and "control") described in this notice? Please explain why.

We agree that the core elements Sensing, Perception, Planning, and Control, are common to current ADS implementations. However, we advise careful design of the framework so that future innovations which might use alternative architectures, for example merged perception and planning, are not precluded.

Question 7. Can you suggest any other core element(s) that NHTSA should consider in developing a safety framework for ADS? Please provide the basis of your suggestion.

In addition to Sensing, Perception, Planning, and Control, two additional core elements should be included: "Localization", and "Prediction". *Localization*, i.e., determining the ego vehicle's position to specified accuracy within its environment need not be fully contained within either Sensing or Perception, and is an element that can be isolated from the others from a validation and testing standpoint. Similarly, *Prediction* of the actions of other vehicles, bicyclists, pedestrians or even animals, is a capability independent of the related elements Sensing, Perception and Planning. The results of Prediction inform Planning. In the medium to long run, Prediction will become an increasingly important contributor to advanced ADS safety.

Independent of the six core elements within the ADS control system, the verified safety-monitor, real-time error handling, and incident reporting telemetry will be active. It is possible these orthogonal functions may also belong on the list of core elements to be assessed within the ADS safety framework.

Question 8. At this early point in the development of ADS, how should NHTSA determine whether regulation is actually needed versus theoretically desirable? Can it be done effectively at this early stage and would it yield a safety outcome outweighing the associated risk of delaying or distorting paths of technological development in ways that might result in forgone safety benefits and/or increased costs?

At the early- to mid-stages of industry ADS development, the technology is a moving target. We should treat the relationship among ADS technology developers and NHTSA as a partnership, advancing toward the shared goal of increased theoretical understanding and practical experience. Information sharing among all parties will be helpful in advancing the ADS state of the art. The industry and academic players at the forefront of ADS development take behavioral safety seriously because it is in their interest to do so. Public-road test and development fleet sizes are small, with hundreds to a few thousand of vehicles throughout the United States. International consensus ADS Behavioral safety standards and guidance, for example ISO 21448 (SOTIF) and UL 4600, are in active discussion and development. Regulation in the form of requiring compliance to industry standards will be appropriate when deployed fleet sizes are large, e.g., in the millions of vehicles, and when consensus has solidified on effective ways to validate and verify behavioral safety.

Voluntary reporting on ADS safety performance under an NCAP model will provide early useful information to the public, and allow NHTSA to develop and iterate metrics that may eventually mature into regulation mandates.

Some states have published guidelines for public road ADS vehicle testing, addressing e.g., whether safety engineers should be present along with safety drivers, what level of driver training is needed, whether collision avoidance systems must be installed. NHTSA may be able to draw from this work to create unified Federal guidance. Pennsylvania's ADS guidance document provides a useful example: https://www.penndot.gov/ProjectAndPrograms/ResearchandTesting/Autonomous%20_Vehicles/Documents/PUB_950_9-20.pdf

Question 9. If NHTSA were to develop standards before an ADS-equipped vehicle or an ADS that the Agency could test is widely available, how could NHTSA validate the appropriateness of its standards? How would such a standard impact future ADS development and design? How would such standards be consistent with NHTSA's legal obligations?

NHTSA could release early draft proposed standards in the form of voluntary guidance open to feedback by industry experts, including those who have already contributed to international consensus safety standards. NHTSA would invite stakeholder and practitioner comment, similar to the association working groups that have created mature industry standards such as ISO 26262. This approach would be useful both for proposed qualitative standards, e.g., guidelines for safety case argumentation and independent evaluation, and for proposed objective standards, e.g., specific curated scenarios to be exercised through hardware-in-the-loop simulation or using physical vehicles.

Ideally, NHTSA should not need to issue its own standards but should identify which industry standards are expected to be adhered to by the manufacturers. NHTSA could require specific interpretations of the international standards (e.g., expected ASILs for specific implementations or specific metrics).

Question 10. Which safety standards would be considered the most effective at improving safety and consumer confidence and should therefore be given priority over other possible standards? What about other administrative mechanisms available to NHTSA?

NHTSA should give priority to the ISO 26262 Functional Safety, ISO 21448 SOTIF, and ISO 21434 cybersecurity industry standards, as these are the product of industry expert working groups and already have consensus. Beyond these, NHTSA should begin the curation of a test scenario simulation suite, initially derived from historical accident archives or scenarios contributed by ADS stakeholders. This suite is important because it will evolve into the primary pre-deployment objective means of demonstrating ADS safety performance over the long run as the industry matures and can form the basis of FMVSS.

To promote consumer confidence, NHTSA should publish summary safety performance data, showing to what extent an ADS vehicle exceeds the minimum standards which it is required to meet, using an NCAP approach.

Question 11. What rule-based and statistical methodologies are best suited for assessing the extent to which an ADS meets the core functions of ADS safety performance? Please explain the basis for your answers. Rule-based assessment involves the definition of a comprehensive set of rules that define precisely what it means to function safely, and which vehicles can be empirically tested against. Statistical approaches track the performance of vehicles over millions of miles of real-world operation

and calculate their probability of safe operation as an extrapolation of their observed frequency of safety violations. If there are other types of methodologies that would be suitable, please identify and discuss them. Please explain the basis for your answers.

Explicit rule-based methodologies are appropriate to establish very basic minimums for ADS Behavioral performance. Rules can be thought of as a transcription or extension of each state's motor vehicle department driver manual (or, for Federal purposes, a composite of the driver manuals of all states). Rule-based assessments are necessary but not sufficient to determine ADS safety performance. In addition to sample vehicle testing in controlled conditions prior to approval, there will be value in ongoing embedded, automated, independent monitoring against a set of rules in development and operational fleets, using rule monitors which run continuously alongside the ADS control system.

Statistical fleet monitoring comprises after-the-fact counting of fatality, collision, incident, passenger-registered discomfort and other anomaly events against miles and/or time driven. Events should be tagged with the "degree of difficulty" of the driving situation (e.g., rural open highway vs. congested chaotic urban rush-hour and weather conditions). Fleet statistics inform long-run safety performance measures and provide a means to quickly notice sudden degradations in safety and take quick action in the form of a "recall-in-place" through temporary ODD restriction, rollback to a prior OTA version, or outright disablement of the ADS function.

A third, important, assessment methodology is a massive, efficient, curated simulation suite of must-pass stress test cases. When statistical fleet monitoring is not yet available, this is the most effective means to demonstrate the safety performance of a not-yet-deployed ADS.

Question 12. What types and quanta of evidence would be necessary for reliable demonstrations of the level of performance achieved for the core elements of ADS safety performance?

Before a large fleet is approved and in operation, the most effective objective way to demonstrate its ADS performance will be via hardware-in-the-loop simulation, where a curated set of millions of well-chosen stress-case scenarios can be exercised during qualification and certification. Physical vehicle test-track and road testing have their place as a tangible way to correlate simulation results to on-road performance but are far more limited than high-fidelity simulation in their ability to probe the space of ADS situations and scenarios, due to cost and time limitations.

Once a fleet is approved and in operation, its own collision, incident, and anomalous event statistics will provide ongoing evidence of ADS safety performance. Reported fleet events will provide information to continuously improve the standard qualification simulation scenario suites curated by the manufacturers and by the Agency.

Question 13. What types and amount of argumentation would be necessary for reliable and persuasive demonstrations of the level of performance achieved for the core functions of ADS safety performance?

Qualitative safety performance argumentation can focus on elements of the ADS control system with verifiable safety attributes, for example the safety or rule monitors which override the primary control system to inhibit unsafe action. Driving policies such as the NVIDIA Safety Force Field (https://www.nvidia.com/content/dam/en-zz/Solutions/self-driving-cars/safety-force-field/an-introduction-to-the-safety-force-field-updated.pdf) can be mathematically proven safe. For the core functions where inductive or deductive absolute proof is not possible, then acceptable failure rates must be derived from system failure rate goals and justified. If demonstrated empirically, the test methodology and results must be made available to assessors.

The Perception core element presents a noteworthy challenge. Perception must be designed and validated to anticipate e.g., weather effects, atmospheric inversions (mirages), optical illusions, glare, noise in LIDAR or radar returns, worn or obscured roadway markings or signage, ambiguous objects, and rare combinations of events. Forming accurate spatiotemporal models of the surroundings requires sophisticated interpretation of sensor outputs, taking account of sensor limitations. Human drivers often compensate for our own imperfect senses. Human strategies, including caution in the face of uncertainty, plausible physics of motion, and semantic situational awareness can inform ADS approaches to this problem.

B. Question About NHTSA Research

Question 14. What additional research would best support the creation of a safety framework? In what sequence should the additional research be conducted and why? What tools are necessary to perform such research?

NHTSA Research should extract the maximum possible insight from every detail available throughout decades of legacy state, county, and municipality accident data history. This will serve two primary purposes: (1) Establish baseline reference expectations for accident data, by severity of accident, organized by degree of difficulty of the driving situation, including congestion, visibility, road conditions, and weather conditions; and (2) Generate specific stress-case scenarios which will become part of the curated simulation qualification suites.

Activities needed to enable this research include the creation of a standardized Scenario Description Representation (SDR), and creation of tools which can import the plain text and diagrams from accident reports archived throughout the U.S. into an SDR.

Ongoing research will be important as ADS technology evolves. Ever-increasing safety performance targets may, for example, require higher levels of ADS situational awareness. NHTSA should fund research and promote and encourage industry / association standards activities to advance analysis techniques to keep pace with these more advanced Artificial Intelligence systems.

NHTSA research into driver interaction with assistance and automation functions, including driver reaction times for driver take-over, and the efficacy of driver monitoring would be useful.

C. Questions About Administrative Mechanisms

Question 15. Discuss the administrative mechanisms described in this notice in terms of how well they meet the selection criteria in this notice.

Performance-oriented, technology neutral, objectively measured administrative mechanisms are best suited to the current early state of ADS maturity. It is preferred to issue proposed, draft, regulation standards in the form of guidance, to invite feedback from research and development experts. The safety-oriented consensus standards efforts have followed a similar path, starting with informative standards which have progressed to normative standards when experience has accumulated.

Question 16. Of the administrative mechanisms described in this notice, which single mechanism or combination of mechanisms would best enable the Agency to carry out its safety mission, and why? If you believe that any of the mechanisms described in this notice should not be considered, please explain why.

Of the mechanisms described, performance-oriented, objective standards, informed by the research described in the answer to question 14, would best enable the agency to carry out its safety mission, early in ADS technology maturity. Over the longer term, all of the mechanisms described are suitable to be considered.

Question 17. Which mechanisms could be implemented in the near term or are the easiest and quickest to implement, and why?

Informative voluntary guidance, including recommendations for e.g., defensive driving policies, would be the easiest to implement, and can open channels for productive dialogue among suppliers and NHTSA. It would be highly valuable for NHTSA to begin to release data which supports the setting of ODD-specific safety metric targets as described in the answer to question 14.

Question 18. Which mechanisms might not be implementable until the mid or long term but might be a logical next step to those mechanisms that could be implemented in the near term, and why?

Mechanisms which presume consensus or homogenous ADS approaches are best to defer until the mid to long term. The logical next steps beyond the near term are the continued expansion in depth and breadth of curated simulation scenarios and the insight extracted from historical incident data. Also, in the medium term, continued advancement of guidance, or the introduction of mandatory standards for fleet incident reporting will be useful to set the stage for the information flow that will inform long-run safety performance advancement.

Question 19. What additional mechanisms should be considered, and why?

NHTSA can draft and invite feedback on FMVSS before ADS technology maturity in certain areas. Some examples:

1. Human-machine-interface (HMI) for the presentation of safety-critical information, and for the controls which call for an emergency stop for passengers in a Level 4 or Level 5 Robo-taxi.
2. During Level 4 or Level 5 degraded operation in response to a malfunction or limitation, set minimum availability and performance goals for the electronic backup, since the human driver is not available to act as backup.
3. HMI for driver notifications in Level 2 or Level 3 vehicles.
4. Rules which define when a manufacturer must issue a recall-in-place (e.g., ODD restriction or OTA version rollback) in response to fleet statistics which indicate a safety regression.

Question 20. What are the pros and cons of incorporating the elements of the framework in new FMVSS or alternative compliance pathways?

FMVSS are prescriptive, and subject to objective evaluation. We have outlined an approach wherein a large suite of simulation verification tests, accumulated as the industry develops experience, forms the basis of a comprehensive FMVSS. But it is premature to rely on FMVSS while ADS technology is still in development. Safety case argumentations and white-box evaluations are more appropriate during early phases.

Question 21. Should NHTSA consider an alternative regulatory path, with a parallel path for compliance verification testing, that could allow for flexible demonstrations of competence with respect to the core functions of ADS safety performance? If so, what are the pros and cons of such alternative regulatory path? What are the pros and cons of an alternative pathway that would allow a vehicle to comply with either applicable FMVSS or with novel demonstrations, or a combination of both, as is appropriate for the vehicle design and its intended operation? Under what authority could such an approach be developed?

For the purpose of this response, FMVSS for ADS will be assumed to include the requirement to pass verification tests that are both test-track and simulation based on curated suites of stress-case test scenarios. These test suites will evolve on a regular basis, for example, annually or quarterly. Off-cadence urgent test suite updates may occur if a critical enough condition is discovered in the field.

There should be a mechanism for manufacturers to raise good-faith challenges as to whether certain scenario(s) properly belong in the test suite. Assuming the test suite is well-curated in this way, it will be reasonable to require compliance with the FMVSS verification testing.

That said, flexible alternative demonstrations of ADS competence can be useful to provide insight into future innovations that might lead to more efficient or effective compliance approaches. This would be considered collaborative research between NHTSA and suppliers, under the authority to issue voluntary guidance, not an alternative means to show compliance.

D. Questions About Statutory Authority

Question 22. Discuss how each element of the framework would interact with NHTSA's rulemaking, enforcement, and other authority under the Vehicle Safety Act.

During early ADS development, some companies intend to use human remote operators who may remotely control the vehicle in the event a development or production driverless vehicle encounters a situation that it is not able to handle while on public roads. The Agency may need to issue rules clarifying whether those remote human drivers must be physically located within U.S. territory, or to assure by legislative action that the activities of those remote drivers are subject to NHTSA authority wherever they are.

Although we believe that cybersecurity will become accepted as one of the pillars of dependable ADS, legislative action or court decisions might be needed in order to clarify NHTSA's authority in the cybersecurity domain, at least as it pertains to ADS integrity and resilience to malicious activity or misuses.

The research into accident characterization and the qualification test curation aspects of the framework fall within NHTSA's authority as expressed in the Vehicle Safety Act.

Question 23. Discuss how each element of the framework would interact with Department of Transportation Rules concerning rulemaking, enforcement, and guidance.

Motor vehicle standards shall be practicable and be stated and testable in objective terms. This motivates our call for simulation demonstration tests and in-field incident reporting to be defined in a universal, supplier-independent manner.

Question 24. If your comment supports the Agency taking actions that you believe may fall outside its existing rulemaking or enforcement authority, please explain your reasons for that belief and describe what additional authority might be needed.

We do not believe any of the recommended actions fall outside of the Agency's existing authority.

Question 25. If you believe that any of the administrative mechanisms described in this Notice falls outside the Agency's existing rulemaking or enforcement authority under the Vehicle Safety Act or Department of Transportation regulations, please explain the reasons for that belief.

Not applicable.