MITRE Corporation Response to the NHTSA ANPRM for the Framework for Automated Driving System Safety

Docket No. NHTSA-2020-0106

Authors:

*Miles Thompson (*miles@mitre.org*)*

*Colin Gladding (*cgladding@mitre.org*)*

*Jessica Lascara (*jlascara@mitre.org*)*

*Erik Phelps (*ephelps@mitre.org*)*

*Fei Sun (*fsun@mitre.org*)*

Contents

Introduction

MITRE's response to this Request for Information provides our recommendations on the design of a formal Automated Driving System (ADS) Safety Framework—one that can scale and evolve as technology advances and that will encourage a culture of safety in all phases of ADS developments. We do not expect our recommendations to be implemented wholly by NHTSA. This challenge will require government, industry, and academia partnerships and MITRE's experience working with them toward the betterment of public safety. Our recommendations:

- address key components needed to objectively define, assess, and manage the safety of ADS performance while ensuring the needed flexibility to enable further innovation.
- are based on our work in developing appropriate safety risk management approaches for multiple industries, including recent work with ADS developers, and our activities assessing and evaluating technologies and solutions for autonomy used in both civilian and military applications.
- focus on the technical aspects of the ADS Safety Framework, and
- provide an approach to transform specific technical recommendations into a comprehensive ADS Safety Framework, which identifies and incorporates all other needed components such as governance, administration mechanisms or statutory authority.

MITRE's mission-driven teams are dedicated to solving problems for a safer world. Through our public-private partnerships and federally funded R&D centers, we work across government and in partnership with industry to tackle challenges to the safety, stability, and well-being of our nation.

As a not-for-profit organization, MITRE works in the public interest across federal, state, and local governments, as well as industry and academia. We bring innovative ideas into existence in areas as varied as artificial intelligence, intuitive data science, quantum information science, health informatics, space security, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resilience.

We operate FFRDCs—federally funded research and development centers. FFRDCs are unique organizations that assist the United States government with scientific research and analysis; development and acquisition; and systems engineering and integration. We also have an independent research program that explores new and expanded uses of technologies to solve our sponsors' problems.

MITRE's transportation expertise spans a wide variety of capabilities, including data and analytics, safety management systems, cyber, autonomy, and spectrum assessment for vehicle to everything (V2X). MITRE also employs experts in safety, the automotive industry, fleet data collection, and autonomous vehicle (AV) development and testing.

We have already proactively engaged with the National Highway Traffic Safety Administration (NHTSA) and the ADS community to share our perspectives on an ADS Safety Framework and welcome the opportunity to continue these conversations in the future.

Goals for ADS Safety Framework

All recommendations contained in this response are tailored to the following MITRE-recommended minimum set of goals for a standard ADS Safety Framework to best serve the public interest. The ADS Safety Framework must:

- Foster continuous improvement of safety in all aspects of automated driving. This means it should be adaptive, incorporating new evidence as it comes to light, while also incentivizing the industry to invest in and share new best practices for safety. This requires the cultivation of a safety culture in the industry, including strong support for independent checks and balances and the empowerment of all employees to observe and participate in the safety process.
- Support iterative design and development processes for new ADS technologies and require technology developers to consider the implications of future progress in safety analysis and engineering.
- Include a mechanism for proactive recognition of issues before they become safety hazards. The design of ADS creates discontinuities in performance, meaning that a systemic increase in near-misses may indicate a catastrophic hazard is on the horizon.
- Establish new expectations for collection and the potential safety benefits of the unprecedented volume and detail of data generated by future ADS-enabled vehicles. Modern vehicles are being fundamentally changed by the demands of ADS features: electronic control modules are multiplying, onboard networks are expanding, and automated vehicle control functions are assuming control in more safety-critical situations. A single commercial ADS can produce and process terabytes of data for each hour of driving, much more than can be stored locally, offloaded quickly, or analyzed in full. Most of this information will be unnecessary for safety analysis, so it is critical to establish a minimal common data standard for information needed to assess safety performance, as well as an approach for evaluating what information is most essential.
- Support safety inspection and certification of ADS-equipped vehicles. Safety inspections are nothing new for human-driven cars, but they are similarly needed for operators of ADS fleets and industry developers.

It is important to note that ADS introduce a paradigm shift. A human-driven vehicle can be as safe as the driver is, and each driver is different. There are as many unique experiences as there are humans. One human driver who encounters a hazardous scenario may learn from that experience and warn others of that scenario. An ADS that encounters a hazardous scenario can inform all other ADS in a relatively short time, leading to exceptionally fast collective safety progress. The NHTSA ADS Safety Framework must prioritize this collective understanding of known hazardous scenarios to leverage this unique advantage of ADS.

Recommendations

The following recommendations address specific components of the advance notice of proposed rulemaking (ANPRM) but are part of a larger, more complete MITRE safety approach and concept.

Summary of Recommendations

This is a list of the recommendations found within this document.

R1. MITRE recommends that NHTSA leverage existing industry Safety Management Systems (SMS) for the federal ADS Safety Framework, including use of MITRE's prior work on SMS, for example.

R2. MITRE recommends that NHTSA adopt a three-tiered safety assessment and management approach, from design to testing to release, with detailed goals and processes in each tier.

R3. MITRE recommends that the NHTSA ADS Safety Framework take a system engineering approach to safety that includes networked support systems, infrastructure, other ADS, and potential human intervention.

R4. MITRE recommends that the NHTSA ADS Safety Framework use the described high-level capabilities of an ADS (sensing, perception, planning, and controls) but also provide a formal definition of each to ensure transparency across the industry as suggested in this response.

R5. MITRE recommends that NHTSA ADS prescribe safety requirements that are dependent on the known hazards within a specific operational design domain.

R6. MITRE recommends that NHTSA adopt a system engineering approach to revising existing federal motor vehicle safety standards (FMVSS), driven by testable performance-based standards derived from implementation-independent safety requirements instead of prescriptive regulation.

R7. MITRE recommends that NHTSA incorporate a continuous validation monitoring effort, operated by the developers, NHTSA, or a third party, to show that the ADS exhibits safe operations as intended, captures anomalous performance, and provides input into the database of known hazardous scenarios as new ones are encountered.

R8. MITRE recommends that statistical analysis of safe driving performance continue to be a part of the ADS Safety Framework. Statistical analysis is not sufficient, but when paired with design risk mitigation, they demonstrate the ADS was both designed to be safe and performs safely.

R9. MITRE recommends NHTSA incorporate a model-based safety analysis approach as part of the tier one risk management to deliver safety evaluations and recommendations for manufacturers throughout their ADS risk management tiers.

R10. MITRE recommends that new guidance and future standards for ADS data logging should be organized to address the exploding volume and complexity of available data and to support a safety methodology that can address new ADS-specific risks and hazards.

R11. MITRE recommends that NHTSA's ADS Safety Framework include a voluntary centralized program of data collection, standardization, analysis and sharing.

R12. MITRE recommends that a common set of known hazardous scenarios be maintained and continuously updated to raise the minimum bar of safety across the entire ADS industry.

R13. MITRE recommends that NHTSA carve out a section of spectrum for V2X (vehicle to everything) communication to better facilitate data sharing between vehicles and data gathering for continuous monitoring, and to protect future infrastructure communication.

R14. MITRE recommends that NHTSA use a form of logic modeling to combine the recommendations from this document with others into a complete ADS Safety Framework.

ADS Safety Framework References

The ADS Safety Framework that NHTSA will implement can pull from many existing resources and industries in order to get the framework quickly assembled. Industry-independent solutions like Safety Management Systems[1] (SMS) and Capability Maturity Model Integration[2] provide examples for iterative framework regulation of complex systems in public sectors. Performance and safety standards—such as ISO/PAS 21448 (upcoming standard on Safety of the Intended Function [SOTIF]), ISO 26262 (functional safety standard used in the automotive industry), UL4600 (standard for performance evaluation of safety cases in autonomous vehicles), and MIL-STD-882E (systems safety standard for the Department of Defense)—provide excellent resources for technical evaluation mechanisms and for assessing technical risk. Other government agencies also offer great examples, like the Federal Aviation Administration, which implemented industry-wide data collection systems to ensure the safety of aircraft and their use.

Because there is a lot of information in those resources, it would be an unwise use of this response to copy directly from them. Instead, this response focuses on recommendations that gradually build on one another to propose important elements that MITRE believes should be included in a successful ADS Safety Framework and that may be overlooked in the current automotive regulations.

Graduated Risk Management with a Three-Tier Approach

Focusing first on technical safety assessment, MITRE recommends that NHTSA adopt a graduated approach with a three-tier approach that gives the best opportunity to ensure the safe operation of ADS on public roads.

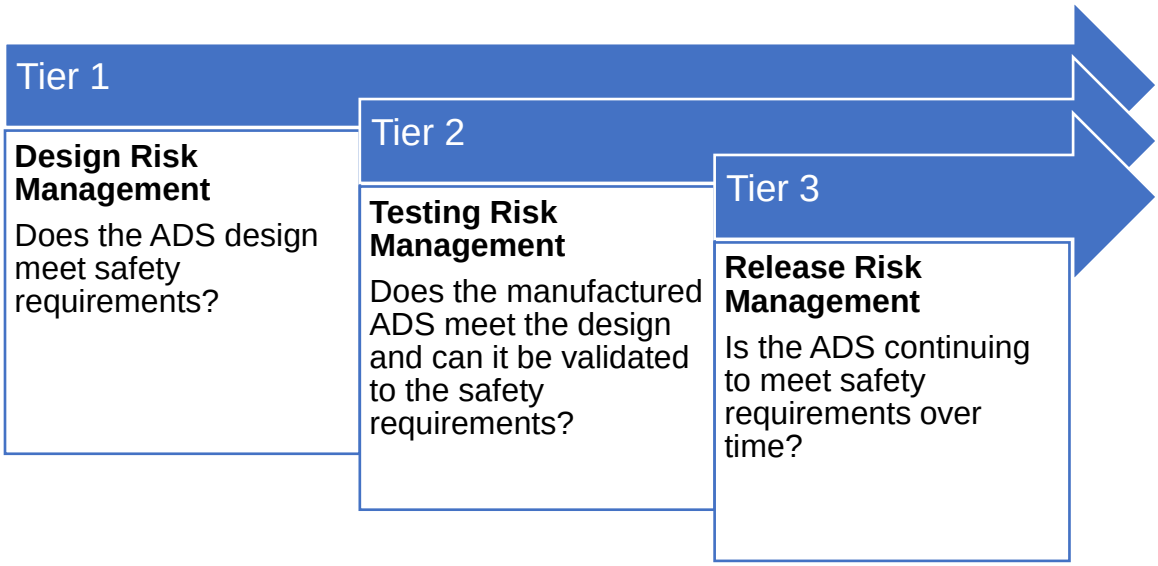

Figure 1: Three-Tier Approach for Time-Limited Certificate to Operate

[1] Safety Management Systems, SMS Explained, at the Federal Aviation Administration (https://www.faa.gov/about/initiatives/sms/explained/)
[2] Capability Maturity Model Integration at the CMMI Institute (https://cmmiinstitute.com/)

Tier 1: Design Risk Management

The first tier, design risk management, is where ADS developers must provide sufficient design evidence to show their ADS will safely handle the set of scenarios that are known to be present within the ADS's operational design domain (ODD). Because there is not necessarily a functioning system available to test in the design phase, MITRE suggests that this design risk management be applied to functional models of the system's safety features. Requiring an ADS developer to demonstrate how the intended development of their system will safely handle the known hazardous scenarios not only promotes public confidence in the product, but also reduces the safety risk due to technical debt, which can be rampant in machine learning (ML) systems. Technical debt is present in a system when a developer makes a system that satisfies the requirements but cannot explain how or why it does. Many perception systems in ADS leverage ML image classification systems, which introduce technical debt by design.

Tier 2: Testing Risk Management

The second tier is integration and data capture risk management on prototype systems. At some point in the development process, the ADS developer will have a functioning system and will request public road access for testing. For that access to be granted, the ADS developer will need to show their prototype system's safety does function in alignment to the original model. This can be done with simulation but must always also include some live tests to check that the data capture mechanism for monitoring the safety of the ADS is operating as expected and integrates with the greater NHTSA data collection effort. Importantly, this testing should also include loss of remote connection, leaving the ODD, and delayed data offload in cases where important data could be lost or overwritten if offload is not timely.

Tier 3: Release Risk Management

The third and final tier, release risk management, is based on continuous validation. When NHTSA has received enough evidence to say an ADS is safe to be released to the public, this statement is based on collected evidence of past performance. Since the environment changes, the vehicle changes (all software needs updates), and laws and regulations change, a method of monitoring known as continuous validation is necessary to monitor the ADS and ensure it continues to operate safely after this initial evidence is collected. Important metrics will need to be periodically gathered and assessed to determine where the safety is increasing, decreasing, or has crossed some unsafe threshold where the ADS needs to revert to a lower tier for more aggressive development updates as a redesign or prototype.

To assist in continuous validation, NHTSA has a new tool at its disposal, the time-limited certificate to operate. Traditional vehicles that meet the FMVSS are released to the public and considered safe until enough evidence is gathered to say otherwise, which usually ends in the developer recalling the unsafe products to avoid regulatory punishment. With continuous validation, the certificate of safe operation can be provided on a time-limited period, based on the quality and frequency of the metric data. If the metric data stops, that may be enough justification for NHTSA to say it can no longer determine the safety of the ADS and it should be removed from public roads.

Safety Management Systems for Holistic Risk Management

In December 2020, MITRE published *Management of Safety Risk in Automated Driving Systems*.[3] In this article, safety experts at MITRE leveraged their decades of experience in

[3] Management of Safety Risk in Automated Driving Systems

different industries to outline an adapted SMS approach for ADS developers. The SMS covers various tiers of the system engineering process, from design and testing/demonstration to vehicle deployment. The safety experts also provided recommendations for safety advancements across the ADS industry through collaborative research, benchmarking, and data sharing.

Although the article focuses on the ADS developer's perspective on how to implement an SMS within their product lifecycle, it also serves as a good reference when agencies such as NHTSA seek to structure a federal framework for Automated Driving System safety. By aligning with industry SMS, NHTSA can leverage existing efforts that original equipment manufacturers (OEMs) made for their ADS safety and utilize the NHTSA ADS Safety Framework for bridging gaps and addressing priorities. This approach also promotes collaboration and therefore eases adoption within the industry.

For example, MITRE's SMS recommends that OEMs incorporate safety features early in ADS design, thus NHTSA should consider developing corresponding standards and methodologies within its framework to evaluate the adequacy of safety designs. Since the SMS helps OEMs develop their evaluation plans based on operational risks and safety criticalities, NHTSA can guide the determination of risks and criticalities and therefore derive the scope, content, and frequency of safety evaluations when assessing OEM demonstrations. In addition, SMS suggests that OEMs implement a confidential employee reporting system for safety concerns that extend beyond the ADS development phase to the deployment and operations phase. Correspondingly, NHTSA should include in its framework a component for collecting inputs from third parties, such as OEMs and the general public, regarding issues and scenarios that were not expected or not included in current evaluation protocols.

The *Management of Safety Risk in Automated Driving Systems* document also discusses the benefits of industry-wide research on developing safety benchmarks and metrics, performance indicators, and data sharing mechanisms and protocols.

Core Elements of ADS Safety Assessment

ADS are complex systems of software and hardware that add new levels of difficulty when assessing their safe operation for even a limited ODD. MITRE suggests taking a system engineering approach in order to decouple the ADS assessment from its implementation and to scope it around holistic safety measures. For example, an ADS that recognizes it has a failure may defer vehicle control to a human driver if available. If not available, that same ADS may defer control to a remote operator to bring the vehicle to a safe state. In both instances, the systems approach provides mitigation for an ADS failure and therefore must be included in the safety assessment.

Focusing inside the vehicle, the ANPRM lists four capabilities to assess: sensing, perception, planning, and control. These four capabilities are the robotics mapping of the OODA loop (observe, orient, decide, act) that is often used to describe how humans ingest their environment. MITRE agrees that these four capabilities are the areas on which to focus the ADS Safety Framework, but these areas are nuanced in ADS. For example, sensing includes not only the vehicle's sensors, but its communications from the operator (e.g., real time maps), its interface with users, and, potentially, V2X communication. All this information is necessary to successfully perceive the vehicle's current state and desired state. The current state, modeled in perception, is often achieved through localization, i.e., the specific type of perception where

the vehicle informs its global position. Although the four capabilities could be broken down into these subcategories, they are better as they are for simplicity and understanding.

To best assess ADS safety, high-level safety requirements must be decomposed and assigned to each capability based on the ODD, object and event detection and response (OEDR), maneuvers, and faults. While the ADS developer is likely the one to perform the decomposition and eventual verification at the capability level, the high-level safety requirements should be developer- and ADS-independent and be maintained by NHTSA. If a developer's ADS is going to drive on the streets of San Francisco without the support of a safety driver, then the ODD (San Francisco) now requires a subset of OEDR and maneuvers for the ADS, similar to a set of use cases, which then trace to safety requirements.

NHTSA has provided these safety requirements in the past in the form of FMVSS. MITRE encourages NHTSA to adopt a system engineering approach to revising existing FMVSS, driven by testable performance-based standards derived from implementation-independent safety specifications instead of prescriptive regulation. MITRE suggests FMVSS continue to apply to ADS as they specify the capabilities that must be within the vehicle to enable the ADS in the appropriate environment.

MITRE also recommends employing a continuous validation monitoring effort to show that the ADS provides performance and safe operations as intended. FMVSS are necessary but insufficient for assessing ADS safety, and they do not provide an assessment of SOTIF performance. Also, it is not enough to be safe; performance must align with the system's expectations to demonstrate sufficient understanding of the ADS.

Statistical measures inform on ADS performance for continuous validation but are not sufficient on their own to determine safety. You also need to do design risk management to ensure that the statistics are indicative of proper, safe performance. Methods like Bayesian change detection, a form of anomaly detection, can provide a rolling window of understanding of the current safety level. Although one might presume more data is better, with ADS that are being updated frequently in an environment that is changing, older data becomes stale very quickly and may no longer reflect the level of safety in the current environment. (However, older data is valuable for baselining and trending purposes, and MITRE recommends retaining a historical archive in some form for those purposes).

While statistical methods can provide some manner of prediction for upcoming safety violations, they are traditionally a lagging indicator and only reveal problems after they have come to fruition versus predicting and preventing issues from occurring. Formal methods and rules-based evaluations provide a standard for assessment in advance and provide a mathematical route for determining the safety of the ADS prior to any on-road performance. Unfortunately, performing these analyses requires intimate ADS knowledge that may not be available to regulators without the developer providing a model of the ADS's safety capabilities. Even if the formal methods proved safe operation, they would prove it only for the model. Regulators would still need statistical performance data to show the model matches the actual ADS.

Model-Based Safety Analysis

MITRE agrees with NHTSA that safety assurance should be integral to ADS design instead of a later consideration in test and demonstration. Therefore, we recommend NHTSA incorporate a model-based safety analysis component in the framework as part of the tier one risk

management to deliver safety evaluations and recommendations for manufacturers throughout their ADS design and development risk management tiers. For example, a manufacturer provides a functional model of its ADS system with all the core elements to the model-based safety analysis component within the framework. The component processes the model, flags the areas within the model that violate safety expectations, and recommends ways of optimization based on the violation. For effectiveness, this process should be automated and thus will require NHTSA to define and develop the following:

- Standards for the ADS model from the manufacturer. This may include the programming language, templates for core elements, and references to data collection.
- The integration process between the model and the model-based safety analysis component within the framework. It is also important to specify communications requirements such as IT and security.
- The key functions within the model-based safety analysis component for automated model processing and optimization with regard to safety guidelines.

Data Logging

On-vehicle ADS data logging will be a key enabler of any program of continuous ADS safety assessment. MITRE recommends that new guidance and future standards for ADS data logging should be organized on the following principles, to address the exploding volume and complexity of available data and to support a safety framework that can address new ADS-specific risks and hazards:

- Safety analysis data requirements should be a central and foremost consideration in future NHTSA activities concerning ADS data logging. The complex and ever-changing nature of ADS systems demands a new data logging approach to support proactive safety analysis and intervention at scale, collecting data to detect safety risks and generate actionable recommendations to avert potential hazards. This includes new logging for new categories of safety hazards that are unique to interconnected software-controlled systems, such as emergent hazards resulting from the behavior of many agents, or hazards resulting from long-term imperceptible shifts in operating conditions.
- NHTSA should encourage standardization for wireless communication to support ADS communications and data collection, including carving out a spectrum of frequencies, allowing ADS data logging monitoring to occur during active use of the ADS. The wireless communication bands have become congested and this will impact both the safety of ADS and the ability of the developers and watchdog organizations to monitor them. Dedicating spectrum will support industry efforts to coalesce on supporting technologies.
- NHTSA should encourage industry to evaluate, adopt, and collaborate on new standardized data logging architectures for hardware, software, and networks to ensure that data critical to accident investigation and safety analysis is captured consistently. Architectures should employ the capabilities of modern computing, communications, and storage technologies to capture data from increasingly complicated onboard software systems and data networks. Data accessibility considerations such as physical access and interfaces and data-at-rest encryption are key and may warrant new or updated common standards.
- Future data logging guidance and requirements should be decoupled from the EDR crash reporting paradigm and its associated goals and assumptions. Current EDR

standards and practices define short-duration logging of a limited set of data elements, triggered by acceleration pulses occurring during crash events. Future ADS systems should use new sensing and perception capabilities to trigger data logging in near-miss events or when expected operating conditions are exceeded, so that these high-risk events can be analyzed. Longer duration logging should capture additional context around crash events as well as metrics describing long-term operating trends that could lead to safety hazards.

- Current EDR standards have a platform-level focus that does not capture the increasingly interconnected nature of ADS-enabled vehicles. Future data logging should extend equal attention to data that describes interacting systems, including the driver state, and the vehicle's physical and digital environments based on perception-based scene understanding and possible future V2X communications.
- NHTSA should work with industry to identify minimal and flexible voluntary standards for onboard data logging of AV-relevant events and data elements, while communicating long-term goals to expand voluntary standards and establish mandatory requirements and standard definitions of events as triggers for data logging, beyond the crash-impulse-based definition.
- Secondary to the requirements of large-scale safety analysis, NHTSA should also consider requirements for ADS data logging to support basic safety inspection of ADS components, to identify potential hazards for individual vehicles. Some categories of hazards, such as those caused by incorrectly calibrated sensors, are difficult for the vehicle to detect but are possible for inspectors to identify.

In addition to the principles above, MITRE recommends that NHTSA's safety framework include a voluntary centralized program of data collection, standardization, analysis, and sharing. An example of a similar systems is the aviation-domain Flight Operational Quality Assurance (FOQA) and Aviation Safety Information Analysis and Sharing (ASIAS) programs. Centralized data collection and sharing through a trusted third party will allow government and industry safety teams to detect rare or emerging issues earlier by analyzing a greater volume of standardized data across manufacturers and models. To adapt these concepts to the ADS domain, MITRE recommends the following actions:

- MITRE recognizes that broad standardization of data formats for on-vehicle logging is unrealistic. A more appropriate solution would be one which requires only that onboard data can be converted into an agreed-upon common representation, such as the FOQA system. This representation might take the form of a common data model into which native data can be translated, or alternatively a set of common performance metrics. Existing efforts to reach consensus around common metrics, such as those by the National Institute of Standards and Technology's ADS-ODD Technical Working Group, may be leveraged.
- Trusted third-party data collection and sharing program.
    - A trusted third-party organization without commercial incentives or conflicts of interest should be chosen to collect, anonymize, and analyze data, and to manage sharing of aggregated data products with researchers and participants. MITRE currently fulfills this role for the Partnership for Analytics Research in Traffic Safety (PARTS), ASIAS, and many other programs across a broad range of domains.
    - Program participants should agree to a common data standard for aggregation across models and fleets, even though native data formats and collection

mechanisms may differ widely. This model follows the example of the FOQA program. This would allow industry flexibility to conform to standards natively, to convert and deliver extracted data themselves, or to work with vendors to do so. In the case of FOQA, multiple vendors have established products to convert raw data frames into a standard format and set of elements, often pairing this service with analysis and processing tools or environments.

  - The trusted third-party organization responsible for data collection and management should work with participants to agree on rules for anonymizing data and aggregated metrics. As the trusted third party for ASIAS, MITRE observes rules for specific data elements to be excluded from all collected data, as well as rules that must be met for any published industry-wide metrics, ensuring operators cannot identify a specific competitor's data.
  - In aviation, privacy concerns for ASIAS were addressed by establishing anonymity of flight data and ensuring legal protections for flight crews and maintenance technicians. A voluntary program offering such protections to participants in the ADS domain who opt in may be a viable option for collecting consumer driving data. Any such option will need to define ownership of logged data and practices for collection, handling and release.
  - In the near term, NHTSA may seek to establish pilot programs for data collection with ADS developers sharing test data and fleet operators sharing on-road data. MITRE recommends starting small, with less sensitive data to prove out the concept. After a successful model is established with those partners and options are explored for consumer data collection, a broader long-term program including consumer on-road data collection would have the greatest impact.

- Data access and analysis
  - NHTSA should be granted access to de-identified data (meaning data is not attributable to the data provider) maintained by the trusted third party, for use in analysis to inform recommendations and policymaking.
  - Participating ADS developers should be granted access to de-identified data for their own analysis and receive results of centralized analysis to support their own safety analysis to identify potential safety or maintenance issues.
  - The public should have access to a set of aggregated and anonymized safety metrics maintained by the trusted third party, similar to current NHTSA databases of EDR crash data, to demonstrate long-term effects of safety improvements and promote public acceptance of AV technology.
- NHTSA should publish guidelines on how data collection could be conducted within an organization as part of a broader safety-oriented culture with involvement from designers, test engineers, worker unions, quality assurance, and other parties.

Common Set of Known Hazardous Scenarios for ADS Assessment

In order to effectively communicate safety expectations to manufacturers, we recommend that NHTSA include in the framework a common set of known hazardous scenarios. The standard UL4600 provides an initial methodology for discovering these hazardous scenarios, and MITRE recommends NHTSA leverage it for this purpose. Additionally, other methods like systems theoretic process analysis can also be used to design a safety case and hazardous scenarios. Through these scenarios, manufacturers would be able to assess their ADS safety not only from the whole system perspective but also at the level of each core element. In the design risk

management tier, developers would be required to show how their design intends to remain safe in each known hazardous scenario. For the integration and prototype risk management tier, a subset of the known hazardous scenarios would be used as test scenarios for the prototype. By passing all the test scenarios in the set, a manufacturer therefore would demonstrate its ability to meet NHTSA's safety expectations.

Providing common, known hazardous scenarios will allow NHTSA to engage with manufacturers throughout the ADS product lifecycle. These scenarios can serve as a reference in ADS design, development, test, and deployment in terms of safety considerations, and they can help set priorities for data collection. Meanwhile, by providing safety guidelines via a common set of known hazardous scenarios, NHTSA assists manufacturers to achieve the same safety goals while maintaining their own innovative solutions.

Further consideration with regard to the common set of known hazardous scenarios may include:

- Standardizing safety test scenarios, including defining metrics for the scenarios such as commonality, criticality, major safety functions under test, major ADS elements under test, etc.
- Defining a process for storing and communicating the scenarios.
- Setting up a known hazardous scenario database (discussed below).
- Voluntary data sharing of new hazardous scenarios as developers encounter them.

Several sections of this response have mentioned a common database of known hazardous scenarios for ADS. Ideally, this would be a NHTSA or community-maintained list of known hazardous scenarios that have been researched, experienced, or predicted, and are indexed in accordance with the Society of Automotive Engineers' (SAE) taxonomy of operational design domain parameters, object and event detection and responses, maneuvers, and faults. The goal would be to have a quick subset of known hazardous scenarios for a new ADS. For example, if an ADS is intended to only operate on U.S. interstates as a temporary chauffeur, then a subset of U.S. interstates as the ODD and attached OEDRs would be selected. This set could then be matrixed with the maneuvers and faults that are ADS-specific for a full list of scenarios to be handled.

Since ADS are still novel technologies, there may not be a large enough repository of data to start this scenario database. Sources such as UL4600 attempt to predict an initial set of hazardous scenarios, but there must be a timely way to add new scenarios as they are encountered, in a way that can be back propagated to other ADS that may already be out of the design phase. ADS developers also likely have large lists of scenarios that caused issues during development and may be willing to provide them voluntarily to seed the initial database. To be as comprehensive as possible, this database should be able to receive scenarios from industry members, local departments of transportation, and even the general public. Additionally, the continuous validation metrics of the release risk management tier would be able to flag crashes, near-misses, disengagements, and other dangerous behavior for automatic review to determine if a new scenario had been encountered.

MITRE recommends the format of scenarios be standardized and defined. One option is the OpenScenario standard, but it is currently too precise in its description of where the actors in the scenario originate and what their behavior will be. Adding a bit of abstraction to OpenScenario

and using the SAE taxonomy to describe the scenarios may be a good way to start organizing them.

Logic Modeling for Successful Outcomes

The recommendations in this document inform on specific components of a larger ADS Safety Framework but do not form a complete concept. MITRE recommends working with NHTSA to start with the desired outcomes of a successful ADS Safety Framework, selecting those recommendations from all responses with value to those outcomes, and building out a complete ADS Safety Framework that supports all the desired outcomes using a logic model. A logic model approach enables organizations to focus investments in areas most important to achieving desired results. The approach begins by articulating the ultimate outcome—a description of a future when a specific problem is resolved or a better state is achieved. The ultimate outcome must focus on those served: customers, constituents, or society at large. From there, interim outcomes—conditions or behaviors necessary and valuable to achieving the end outcome—are defined. Together, the interim outcomes must be sufficient to achieve the ultimate outcome. The ultimate and interim outcomes provide a beacon for aligning programs and developing performance measures. The logic model approach provides a framework for optimizing the return on investments intended to accomplish desired objectives.

Direct Response to Select Questions

This section answers the questions asked in the ANPRM by taking content from the rest of the document. While this section directly answers the questions, better context and rationale is provided in the earlier sections of this response.

Safety Framework

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."

See details in *Summary of Recommendations*.

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?"

As recommended in *R1*, MITRE recommends that NHTSA leverage existing industry SMS, where OEMs develop their evaluation plans based on operational risks and safety criticalities. NHTSA can guide the determination of risks and criticalities and prioritize its monitoring and safety oversight efforts accordingly. More information can be found in *Safety Management* Systems for Holistic Risk Management.

As recommended in *R8* and *R10*, statistical analysis and data logging should also be priorities for NHTSA's monitoring and safety oversight efforts. More information can be found in *Core Elements of ADS Safety Assessment* and *Data Logging*.

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

As recommended in *R4*, *R9*, and *R12*, NHTSA should ensure transparency on core element capability definitions across industry, incorporate a model-based safety analysis component, and maintain a common set of known hazardous scenarios. Statistical analysis and data logging continue to play important roles in effective assessment of core element safety.

More information can be found in *Core Elements of ADS Safety Assessment*, *Model-Based Safety Analysis*, *Common Set of Known Hazardous Scenarios for ADS Assessment*, and *Data Logging*.

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?"

As recommended in *R6*, NHTSA should adopt a system engineering approach driven by testable performance-based standards derived from implementation-independent safety requirements instead of prescriptive regulation.

More information can be found in *Core Elements of ADS Safety Assessment*.

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?"

In addition to recommendations for Question 3, where NHTSA should incorporate a model-based safety analysis component and maintain a common set of known hazardous scenarios, MITRE also recommends that NHTSA adopt a three-tiered safety assessment and management

approach (*R2*) and take a system engineering approach to safety to include networked support systems, infrastructure, other ADS, and potential human intervention (*R3*). Statistical analysis and data logging continue to play important roles in effective assessment of core element safety.

More information can be found in *Graduated Risk Management with a Three-Tier* , *Core Elements of ADS Safety Assessment*, *Model-Based Safety Analysis*, *Common Set of Known Hazardous Scenarios for ADS Assessment*, and *Data Logging*.

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."

As recommended in *R4* and *R13*, MITRE agrees that the core elements of an ADS include sensing, perception, planning, and controls. In addition, NHTSA should carve out a section of spectrum for V2X communication.

More information can be found in *Core Elements of ADS Safety Assessment* and *Data Logging*.

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."

Similar to recommendations for Question 6 (*R4* and *R13*), NHTSA should consider V2X as a core element.

More information can be found in *Core Elements of ADS Safety Assessment* and *Data Logging*.

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?"

No recommendations for this question.

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?"

As recommended in *R2*, NHTSA should adopt a three-tiered safety assessment and management approach, from design to testing to release, with detailed goals and processes in each tier. More information can be found in *Graduated Risk Management with a Three-Tier* .

No recommendations with regard to legal obligations.

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

As recommended in *R1*, NHTSA should leverage existing industry Safety Management Systems (SMS) for the federal framework, including use of MITRE's prior work on SMS, for example. More information can be found in *Safety Management* Systems for Holistic Risk Management and

ADS Safety Framework References.

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."

As recommended in *R8*, statistical analysis paired with design risk mitigation shows that ADS was both designed to be safe and demonstrates safety. More information can be found in *Core Elements of ADS Safety Assessment*.

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?"

As recommended in *R12*, a common set of known hazardous scenarios should be maintained and continuously updated to raise the minimum bar of safety across the entire ADS industry. Evidence of successfully handling those scenarios would be necessary for reliable demonstrations of the level of performance achieved for the core elements of ADS safety performance.

In addition, MITRE recommends in *R7* that NHTSA require a continuous validation monitoring effort, operated by the developers, NHTSA, or a third party, to show that the ADS exhibits safe operations as intended, captures anomalous performance, and provides input into the database of known hazardous scenarios as new ones are encountered.

More information can be found in *Common Set of Known Hazardous Scenarios for ADS Assessment*.

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?"

No recommendations for this question.

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?"

No recommendations for this question.

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."

No recommendations for this question.

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."

As recommended in *R11*, NHTSA's safety framework should include a voluntary centralized program of data collection, standardization, analysis, and sharing. More information can be found in *Data Logging*.

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

As discussed in *Data Logging*: In the near-term, NHTSA may seek to establish pilot programs for data collection with ADS developers sharing test data and fleet operators sharing on-road data. MITRE recommends starting small, with less sensitive data to prove out the concept.

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?"

As a continued discussion in *Data Logging*: After a successful model is established with those partners and options are explored for consumer data collection, a broader long-term program including consumer on-road data collection would have the greatest impact.

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

As recommended in *R13*, NHTSA should carve out a section of spectrum for vehicle to everything (V2X) communication to better facilitate data sharing between vehicles and data gathering for continuous monitoring, and to protect future infrastructure communication. More information can be found in *Data Logging*.

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

No recommendations for this question.

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?"

No recommendations for this question.

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."

No recommendations for this question.

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

No recommendations for this question.

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."

No recommendations for this question.

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."

No recommendations for this question.