WORLD ECONOMIC FORUM

Safe Drive Initiative: SafeDI scenario-based AV policy framework – an overview for policy-makers

WHITE PAPER

NOVEMBER 2020

Cover: Getty Images/Oonal

Inside: Getty Images/IGphotography

Contents

Preface

Michelle Avary,
Head of Automotive
and Autonomous Mobility,
World Economic Forum

Tim Dawkins,
Project Lead, Automotive
and Autonomous Mobility,
World Economic Forum

Measuring and defining safety is one of the greatest challenges in the autonomous vehicle (AV) sector.

AVs promise safer, more sustainable future mobility, but in order for this to be realized, policy-makers and members of the public will need to be assured that the vehicles are safe to operate on the road. Meanwhile, the question of how to define safety for AVs remains a heavily debated subject among industry stakeholders and policy-makers.

One of the challenges faced by policy-makers in considering how to evaluate the safety of AVs is the diversity of the applications of AVs being developed. Different choices in hardware, software and design ethos result in very different safety cases from different AV providers. Even the definition of an edge case can be dependent on the design choices of the AV's perception system. With such variation, how can regulators be assured that two vehicles from different providers, operating on the same street, are upholding the same level of safety?

Hence, through the Safe Drive Initiative, with the support of our industry, government and academic partners, we set out to develop guidance and tools to create a useful, practicable governance approach for AVs – based not on how an AV is built but on how it behaves in the context of its operating environment. Fundamentally, this is predicated on the hypothesis that operational safety for an AV is defined by the nature of the environment in which it is deployed.

In this framework, we provide an end-to-end process through which a policy-maker can set a vision for the future role of AVs in society and develop a safety assurance process tailored to the specific challenges of the streets and highways on which they are deployed. This approach enables open innovation because it does not prescribe specific solutions from AV developers; it is also vehicle- and solution-agnostic. Moreover, we propose an approach that builds upon existing standards and safety research where possible to allow for the harmonization of common elements.

One of the fundamental success factors in deploying this framework is continued direct engagement by the regulator with the AV industry throughout the process. Not only does this enable an agile, performance-based regulatory structure to be created but it should also deepen the regulator's knowledge of this technology sector and establish lasting trust with AV developers.

As soon as this framework has been published, we are seeking to collaborate with one of our government partners to pilot this approach and refine our guidance to further develop an implementation toolkit to enable other regulators to scale and adopt the framework.

1 Executive summary

The Safe Drive Initiative seeks to bridge the gap between the industry's expertise in AV safety and the regulators' desire to set policy that safeguards AV deployment. The technical complexity of AV systems has presented a challenge for regulators, who seek to assure the public of AV safety but are not sure how to do so. An effective safety assurance process must strike a balance between ensuring safety for all road users and while not making unreasonable demands on AV providers.

This initiative proposes a data-driven, scenario-based assessment using a graduated approach that will enable regulators to independently assess the safety of an AV system. It is data-driven in the sense that the inputs and outputs to the process are concrete and based on real-world information about AV performance within the regulator's jurisdiction. The framework depends on demonstrating safe outcomes in selected scenarios as a common standard of behaviour against which multiple AV platforms can be compared. In addition, this approach tests an AV's performance in increasingly challenging environments, allowing for phased deployment by expanding the available testing area as the AV demonstrates proficiency at each step.

The regulator's primary role in this process is to convene a centre of excellence (CoE) and define the end goal for AVs, in order to set objectives for the policy-making phase. The CoE should have the requisite technical expertise to define the key milestones, measures and specific tests (e.g. on-road testing, controlled environment testing and simulation) to evaluate AV performance from multiple dimensions. These details should align as much as possible to AV developers' existing validation processes.

This framework was developed using the Forum's multistakeholder methodology, by convening technical experts, regulators and industry leaders from around the world to explore and share knowledge and aggregate best practices on AV safety. Reflecting this diverse representation, we recognize that every city, state and country offers different challenges for AV systems and developers. As such, the proposed assessment process is designed to be customized for each jurisdiction's context to bring out the best possible outcomes for all stakeholders.

2 Key terminology

2.1 General terms

- **Automated driving system:** The hardware and software that are collectively capable of performing the entire dynamic driving task on a sustained basis, regardless of whether it is limited to a specific operational design domain.
- **Autonomous vehicle (AV):** A vehicle equipped with an automated driving system designed to function without a human driver as a Level 4 or 5 system under SAE J3016.
- **Dynamic driving task:** All of the real-time operational and tactical functions required to operate a vehicle in on-road traffic, excluding the strategic functions such as trip scheduling and selection of destinations and waypoints.
- **Operational design domain (ODD):** A description of the specific operating domain(s) in which an automated driving system is designed to properly operate, including but not limited to roadway types, speed range, environmental conditions (weather, daytime/night-time etc.) and other domain constraints.

2.2 Scenarios and related terms

- **Scenario:** A traffic situation within the vehicle's operational design domain.
- **Behavioural competency:** A manoeuvre or function that an automated vehicle can demonstrate in various scenarios – for example: Turn Left or Emergency Stop.
- **Scenario-based assessment:** Evaluating a system based on its performance when exposed to a variety of predefined scenarios that correspond to its intended deployment ODD.
- **Minimal risk condition:** A condition of the autonomous vehicle or system to which either the user (safety operator) or the system itself brings the vehicle to reduce the risk of a crash when a given trip cannot or should not be completed. For example, a minimal risk condition might entail "bringing the vehicle to a stop in its current travel path" or "a more extensive manoeuvre designed to remove the vehicle from an active lane of traffic".

In this framework, we refer to three categories of scenario:

- **Qualitative scenario:** An abstract description of a scenario in natural language, with definitions of the traffic situation, driving environment, other vehicles and road users, and environmental conditions – e.g. vehicle in traffic on a three-lane roadway on a summer's day.
- **Logical scenario:** A qualitative scenario that has been parameterized, including possible value ranges for each parameter; it may also include probability distributions for certain parameters – e.g. lane width 2.3–3.5 metres (m), traffic speed 0–30 kilometres/hour (km/hr), temperature 10–40°C.
- **Concrete scenario:** A logical scenario with specified values for each parameter. Such a description is grounded in its environment (context, with its parameters) and includes ego vehicle goals – e.g. lane width 2.3 m, traffic speed 30 km/hr, temperature 23°C.

3 Overview

The Safe Drive Initiative seeks to establish a high-level framework (the SafeDI framework) to enable a regulator and AV developer to collaboratively demonstrate an AV system's capability to operate without intervention from a safety driver. This framework describes a scenario-based operational safety assessment developed through extensive engagement with multiple stakeholders from the AV industry, academia, civil society and global regulators. The guidance is meant to enable regulators and AV developers to develop a shared understanding of the core knowledge required to implement a validation process.

The approach builds on the explicit assumption that an AV's safety can be assessed in the context of its deployment ODD by demonstrating behavioural competence in a range of scenarios. This assessment is intended to be indicative, for the purposes of enabling scaling of AV trials and pilots, rather than an empirical statement of 100% safe operation in all conditions.

While this framework cannot provide a universally prescriptive approach, it does offer high-level guidance on a safety assessment process that can inform collaborations between local regulators and AV developers. Moreover, in the spirit of its organizational mission to improve the state of the world, the World Economic Forum will continue to advocate for and encourage AV operators to share knowledge and learnings on AV safety to benefit the broader industry and the public.

This paper introduces the key elements of the approach and is the second in a series of documents provided by the Safe Drive Initiative. This section describes the guiding principles and approach to the framework, and the next section provides a high-level overview of the framework. Additional technical guidance to support the implementation of this approach is provided in a separate publication.

3.1 Application – who is this framework for?

This framework is intended to be adopted by a regulator or government entity that is responsible for managing AV development and deployment but which does not currently have existing policies that provide it with an oversight or approval process to allow AVs to operate in public. This entity may be a national, state, regional or municipal authority, according to the market, or possibly a coalition of various levels of government.

This framework offers a "fast follower" regulator the opportunity to create a graduated approval process that enables AV deployment, while allowing for a policy-maker to specify its own expectations for demonstrating safe operation in a given environment.

Moreover, it is hoped that the safety assurance approach proposed in this paper provides industry stakeholders with guidance on how to provide clarity to regulators on framing their safety case to a regulator in terms of the operating environment.

No safety assessment framework can avoid the tension between AV providers' need to protect their intellectual property and a regulator's desire to make information public for the purposes of independent assessment and validation. Striking an appropriate balance between these needs will be a challenge, requiring holistic consideration of the broader environment that will be affected by any regulations defining the safe operation of AVs.

3.2 Assumptions

This approach is most useful under the following conditions:

1. A regulator wishes to implement a safety assurance approach that does not rely on industry to self-certify its approach to safety
2. An AV provider has not conducted extensive testing in the deployment ODD within the regulator's jurisdiction
3. The regulator is the sole party responsible for AV governance in its respective market or country or can convene necessary stakeholders to implement such a policy
4. The regulator can assure AV companies that their intellectual property will be protected

While this policy exercise assumes a blank canvas from a policy perspective, markets that have implemented guidance can also benefit from considering this approach. Local nuances and existing policy should also be considered before embarking on this endeavour.

In addition, this framework assumes that the regulator has established a vision or mission statement for the future of mobility in its jurisdiction and the role that AVs will play in this future. Based on our prior AV policy review, the Forum views such a mission statement to be a key success factor in establishing a cohesive policy environment for AVs. In this case, the Forum recommends that such a vision statement should focus on the societal need that can be addressed by AVs and remain technology-neutral.

3.3 Guiding principles of the SafeDI framework

The Safe Drive Initiative seeks to combine the methodologies of ongoing AV safety validation and regulatory initiatives to create an operational safety assessment that is valuable to all stakeholders (Figure 1). As such, the initiative is informed by the following guiding principles:

For regulators and AV developers: Both key stakeholders should benefit from the framework and finds its guidance practical and implementable.

Evaluation uses scenario-based assessment: Key scenarios within the deployment ODD will be identified and used to evaluate the AV's abilities.

Uses a set of common metrics for testing: Commonly accepted metrics should be used to evaluate AV performance (e.g. ODD excursions, crash statistics).

Covers simulations, driving in controlled environments and naturalistic on-road driving: The assessment framework should apply all tools available across the test environments to evaluate AV safety in the most rigorous and efficient ways possible.

FIGURE 1 **The Safe Drive Initiative seeks to combine the best practices from ongoing validation initiatives and regulatory policies. Singapore's Operational Safety Assessment is the only other example of this combination to date (and is not affiliated with the Safe Drive Initiative)**

Existing processes are **challenging for regulators to operationalize**

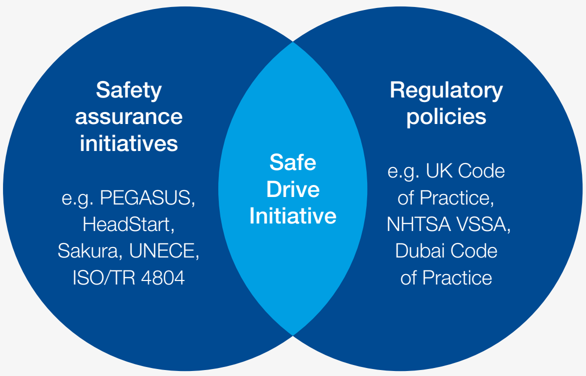

With very few exceptions (e.g. Singapore), **regulatory policies have lacked independent assessment of AV safety**

Existing policies **provide high-level guidance** or **allow self-certification**

The Safe Drive Initiative is a **policy framework that guides regulators through the process of creating a safety assurance programme** that evaluates AV safety

Source:
World Economic Forum / McKinsey & Co. Analysis

3.4 Approach summary

The policy framework comprises four phases (Table 1) that provide a regulator with the necessary knowledge to implement an operational safety assessment of an AV. The process does not require significant technical knowledge to understand it. Instead, it advocates the convening of a group of technical experts who define the specific details of the safety assurance programme in collaboration with regulators and AV developers. However, regulators should still oversee the end-to-end process, ensuring that each step effectively evaluates AV platform safety while minimizing risk to the public.

TABLE 1 **Four phases of the Safe Drive Initiative's policy framework, which can be repeated as necessary to refine details throughout the process**

Prepare	Define	Measure	Execute
Create a centre of excellence (CoE) Convene necessary stakeholders, define the end goal and engage with industry	Define the test ODD Establish the required behavioural competencies for the AV, define the geographic areas and parameters for each interim milestone	Specify on-road, controlled environment and simulation tests, and determine success/ advancement criteria	Conduct tests and collect data from AV providers as necessary, improving the safety assurance process as needed

The first step in the process is to create a centre of excellence (CoE) with the goal of convening the necessary stakeholders, including local, regional and national regulators, academic researchers, industry representatives and municipal leaders. The members of this group will jointly establish the desired end state (i.e. define the deployment ODD) and are to be consulted in the development of the required competencies for AV systems – both how AVs should behave and how AV developers need to demonstrate safe performance.

Via the CoE, regulators and AV providers will then work together to define a series of interim milestones, establish a set of qualitative scenarios and create agreed-upon tests of the AV's competency. These tests first open up small areas for on-road testing, slowly increasing the test ODD to match the desired deployment ODD when the AV's competency to operate safely across defined scenarios has been sufficiently proven. The CoE should also determine metrics for the AV system's successful performance and criteria for advancing the system to the next milestone. Ideally, these metrics and functional requirements should be harmonized with existing standards or governance tools where possible.

Once the process is established, the CoE should oversee the testing and reporting process and seek to continuously improve the validation process.

3.5 Roles

The rapidly evolving nature of AV technology development, coupled with the regulatory structures that are geared towards human-driven vehicles, underscores the need for a multistakeholder approach to AV policy development. This framework is intended to lay the groundwork for a path to navigate the tension between developers wishing to advance technology as rapidly as possible through a transparent assessment process, and regulators wanting to ensure that the technology is sufficiently safe for public operation. The anticipated roles of regulators and AV developers in this process are detailed in Table 2.

An additional success factor in enabling this process is gathering a group of independent experts who can ensure that regulators understand the highly technical documents submitted by developers and translate regulators' wishes into implementable tests by AV developers. These experts can include academics and representatives from industry bodies who can provide objective input throughout the process.

Additionally, the AV developers wishing to operate within the jurisdiction should be directly engaged by the CoE to understand their desires for testing and deployment.

TABLE 2 **Stakeholder roles in the CoE**

Stakeholder	Government	AV developer	Independent technology expert
Typical profiles	Regulator responsible for AV policy (e.g. Ministry of Transport) Law enforcement representative (e.g. traffic police)	Policy or government relations representative Chief safety officer	Technical expert from industry on specific AV components or systems Academic researcher on AV validation
Roles	Define geographic constraints and operating parameters for each milestone's ODD Establish CoE Define qualitative scenario library Establish test success criteria	Provide input on desired ODD Parameterize qualitative scenario library into logical scenarios Prioritize scenario along notable parameters (e.g. routine, critical, edge case and long tail) Suggest tests to demonstrate proficiency in each scenario Demonstrate ongoing safety system (e.g. safety driver, remote safety driver, chase vehicle) Share data from ongoing testing	Review test documentation (e.g. development process, safety record etc.) Objectively mediate process – ensure regulator understands technical documents Assist in translating regulator's desires to actionable requests for AV providers, ensuring that AV providers share accurate test documentation

3.6 How to use this document

For each step of the process, we provide high-level guidance in the form of key steps, guiding principles and suggested best practices and standards where applicable. Each step can be customized to match the needs of the deployment environment. As such, regulators need not strictly follow the steps or guiding principles as laid out, should they not be relevant to a particular context. In particular, jurisdictions that already have AV policy and/or testing in place should adapt their history and experience to this model to create a safety assurance process that suits the particular needs and desires of their communities. This model provides a starting point and broad guidance from which to implement a policy framework that, ideally, all stakeholders will support, from AV developers to members of the public.

4 Process

The approach defined in the SafeDI assessment framework is broadly divided into four phases:

1. Prepare: Convene the necessary stakeholders, define the end goal and fill out the details of the process

2. Define: Establish the required behavioural competencies for the AV, define the geographic areas and parameters for each interim milestone

3. Measure: Specify on-road, controlled environment and/or simulation tests, and determine the success/advancement criteria

4. Execute: Conduct tests and collect data from AV providers as necessary, improving the safety assurance process as needed

The key activities for each of these phases are summarized in Table 3 and presented in more detail in the sections that follow. Additional technical guidance is forthcoming in a future Safe Drive Initiative document to be published by the Forum.

TABLE 3 **Key elements of the SafeDI safety assurance approach to approve an AV platform to operate within a defined ODD**

	Prepare	Define	Measure	Execute
Description	Convene necessary stakeholders, define the end goal and fill out the details of the process	Establish the required behavioural competencies for the AV, define geographic areas and parameters for each interim milestone	Specify on-road, controlled environment and simulation tests, and determine success/advancement criteria	Conduct tests and collect data from AV providers as necessary, improving the safety assurance process as needed
Key activities	Create a **centre of excellence** (CoE) Convene **regulators, AV providers, academics and municipal leaders** to establish process and conduct tests Determine **deployment ODD parameters** to approve AVs to drive in	Determine **key behavioural competencies** for AV platform to demonstrate Define **interim milestones** and **determine subset of ODD for each** (constraining geography, weather, speed limits etc.) Determine **scenario to prioritize at each stage** in the validation process (e.g. critical scenarios earlier, edge cases later)	Identify/generate **scenario library** based on knowledge of ODD and/or naturalistic data Choose **which type of test** (e.g. simulation, controlled track or on-road in ODD) to conduct for each scenario Determine **how to measure the success** of each test, and **when to advance AV systems** to the next milestone Verify **one or more independent mechanisms to handle emergency takeovers** (e.g. safety driver, remote operator, chase vehicle etc.)	Conduct **on-road tests and evaluate simulation results** Collect **data on scenario exposure** and surroundings to refine evaluation

4.1 Prepare – own the process and map and engage relevant stakeholders

Before embarking on the creation of such a policy framework, it is necessary to designate ownership of the process, to ensure that a responsible party is nominated to be accountable for the scoping, implementation and ongoing monitoring and reporting of the policy in place. Hence, this framework strongly recommends creating an AV centre of excellence (CoE) to assume ownership and responsibility for this AV framework, and to become a bridge between the regulating body and the industry.

Profile of a centre of excellence

In many markets, the regulating body responsible for the oversight of autonomous vehicles is likely to be a ministry or department of transport, a department of motor vehicles or a driver and vehicle licensing agency. Naturally, it makes sense that the AV CoE comes from within this organization to complement and strengthen that oversight.

At a bare minimum, the CoE should be a designated individual within this governing body responsible for the AV sector in the market at large. If more resources are available, a broader team of specialists could comprise this CoE, or even a whole department. The CoE should have the combined authority and technical ability to implement the required test procedures.

Ideally, the CoE would have a multistakeholder composition to enable it to holistically consider the challenges of AVs while making the most of available subject-matter expertise. Hence, a joint CoE established in partnership with a university or other research institution would be the ideal recommended profile. A leading example of this approach is Singapore's Centre of Excellence for Testing & Research of Autonomous Vehicles – NTU (CETRAN).[1]

Engaging with key stakeholders

Having designated a CoE, the next step is for the CoE to convene other relevant stakeholders, such as other government bodies, AV providers, academics and municipal leaders that will be impacted by the testing and development of AVs. The purpose of this initial activity is to understand and map these stakeholders' needs and establish a series of end objectives for the testing and development of AVs in the market. An example of this end objective could be: *"Create a testing and deployment programme that enables safe driverless operation of AVs on the streets of City X."*

Having established an end objective, the CoE should prepare for a series of further consultations with AV developers to determine their goals and parameters for their eventual deployment ODD.

Broadly, the deployment ODD represents the largest geographic area in which the AV developers wish to operate, with the broadest set of parameters for weather, time of day and other operating conditions. The deployment ODD should be formally expressed with an agreed ODD taxonomy (e.g. NHTSA[2] or BSI[3]).

4.2 Define – detail interim milestones in a graduated approach

Once the end state is defined in terms of a full deployment ODD, the CoE must develop a series of interim milestones and test criteria that, if successfully followed and satisfied, would provide assurance that an AV provider can operate safely across the deployment ODD. These milestones should be structured to reflect increasing levels of operational complexity for the AV system.

Express interim milestones as a function of the deployment ODD

The purpose of the interim milestones is to provide incremental levels of assurance that an AV is able to operate at an acceptable level of risk. As the environment becomes more complex, the AV should be expected to demonstrate increased maturity to handle critical situations, such as interacting with vulnerable road users and other vehicles.

In this framework, we propose that the interim milestones be set in terms of an incremental and representative series of ODDs. This means working backwards from the deployment ODD to set up a series of ODDs that are an indicative assessment of the roads, scenarios and traffic environment of the next milestone.

For example, the first milestone could be a geographically limited, low-risk area such as a business park available only for testing during off-peak hours. As the AV system demonstrates basic competency, regulators could then allow testing in a wider area with high-quality infrastructure that is more representative of the deployment ODD, but that also limits interactions with vulnerable road users. The testing area for each milestone should continually expand to match the deployment ODD.

FIGURE 2 **Illustrative graduated milestones**

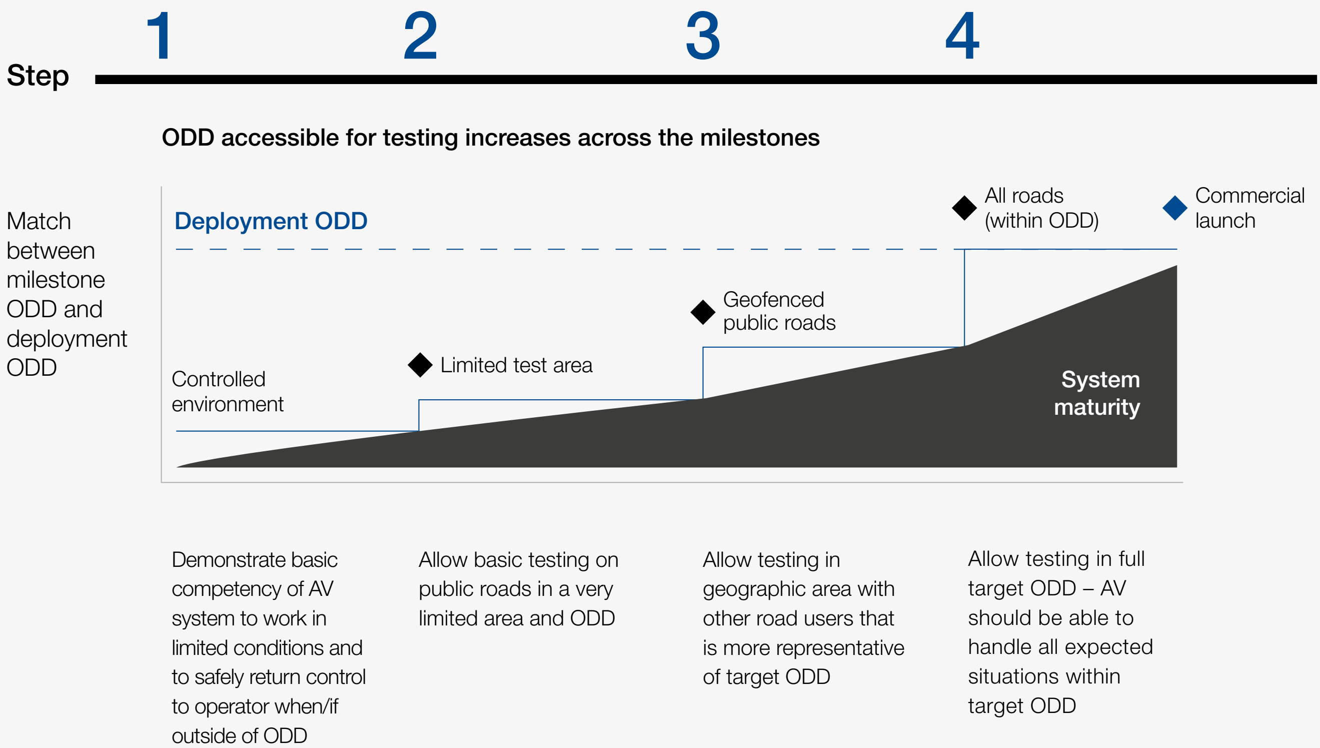

Source: World Economic Forum / McKinsey & Co. Analysis

Specify qualitative scenarios for the interim milestones

Having determined the outcome of each milestone, it is now possible for the CoE to determine the assessment for each milestone.

In this framework, we propose a scenario-based approach to assessment in which the CoE specifies high-level behavioural competencies or qualitative scenarios that are representative of the situations to which the AV will be exposed in each milestone's ODD. The AV developers will be responsible for parameterizing these scenarios into logical and concrete scenarios to create test cases relevant to their vehicle's specific hardware and software stack.

In this approach, a scenario can be considered as the requirement "to perform a manoeuvre, in a specific environment, in certain conditions". Hence, the CoE should view this task as expressed in terms of "what are the behaviours that I want to see demonstrated to know that this vehicle can safely operate in this area?" These can then be translated to scenarios by considering the places or types of driving environment in which one would expect such behaviours to be demonstrated.

The CoE should aim to express these in terms of behavioural competencies or qualitive scenarios, using an agreed scenario description language aligned with the ODD taxonomy. In general, there are two approaches to specifying these initial requirements:

– **Logical approach:** Building on existing sets of competencies (e.g. NHTSA, Waymo) and knowledge of the ODD (e.g. analysing types of road users, intersections) to determine necessary competencies

– **Data-driven approach:** Application of collision data, insurance data etc. to understand key risk areas and extract the necessary behavioural competencies.

Regulators may choose different levels of detail based on their capacity and the unique requirements of their jurisdiction relative to other areas in which the AV platform may previously have been tested.

FIGURE 3 **Life cycle of a scenario from behavioural competency to test case**

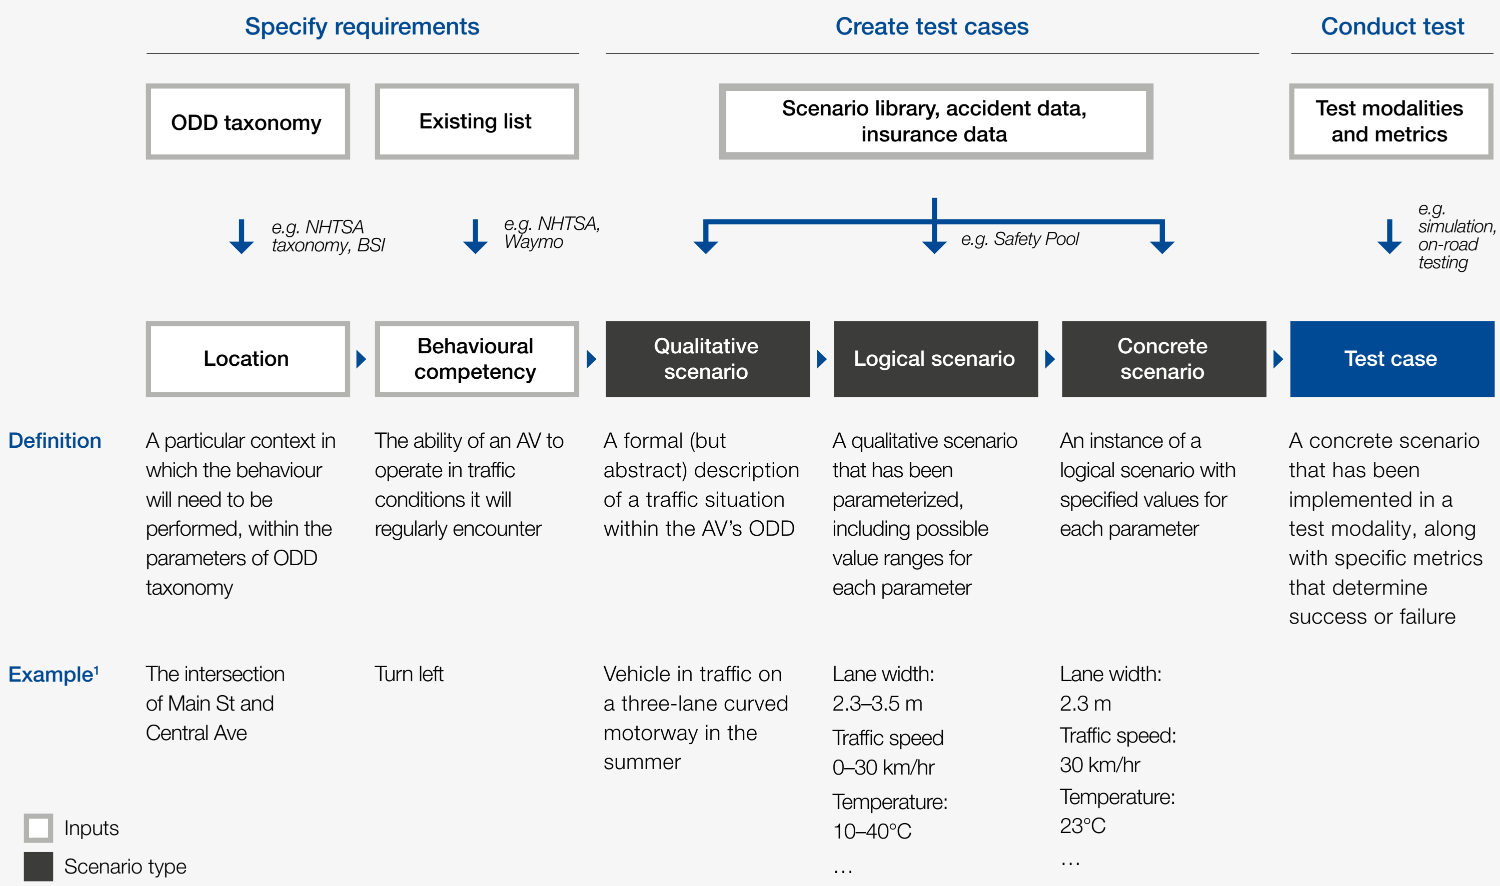

[1]Example based on Pegasus Project https://www.pegasusprojekt.de/files/tmpl/PDF-Symposium/04_Scenario-Description.pdf (link as of 29/9/20).

Source: World Economic Forum / McKinsey & Co. Analysis

4.3 Measure – determine required tests for each milestone

Establish scenario library

Having built a list of scenarios for each ODD, the regulator should now seek to generate, source or otherwise establish a detailed reference dataset of example scenarios that can be used for assessment at each milestone. Additionally, by tagging the scenarios with metadata to link specific elements of the ODD, the regulator can create a searchable dataset based on key parameters.

According to the CoE's preference, the broader policy environment and/or market status quo, the CoE could choose to undertake this development on its own, adopt an existing library, outsource the development or require AV developers to demonstrate compliance. These approaches are summarized in Table 4.

TABLE 4 **Approaches to generating a scenario library**

	High touch	Medium touch	Low touch
Description	Regulator develops own scenario library, either directly or through crowdsourcing of data	Regulator builds on existing scenario library to define competencies	AV developers provide scenario library/competencies they have used to demonstrate performance
Benefits/challenges	+ Data set is drawn from deployment environment and hence fully representative – Expensive, time-consuming, may have limited scenario coverage	+ Reduced cost, builds on existing initiatives for scenario-based assessment, broader scenario coverage – Unlikely to be significantly representative, suitable only for broad demonstration of competence – Requires customization following preset taxonomy	+ Lowest cost and effort to CoE + Generally preferred path for AV developers – Places onus on AV companies to provide safety assurance, reduces insights to CoE and regulator – Less independent verification possible of AV performance within specific jurisdiction

As part of the Safe Drive Initiative, the World Economic Forum has partnered with Deepen.ai and our Automotive and Autonomous Mobility community to create Safety Pool,[4] an incentive-based scenario data exchange among AV developers, which is also intended to function as a reference library for scenario-based assessment. The Safety Pool includes a variety of qualitative, logical and concrete scenarios.

For each milestone, the CoE should prioritize scenarios the AV developers must demonstrate to provide a suitable level of safety assurance prior to deployment in the corresponding ODD. The scenarios can be grouped into four categories, as expressed in Figure 4.

FIGURE 4 **Regulators should qualitatively define scenarios, and prioritize which parameters within each scenario to test at each test milestone**

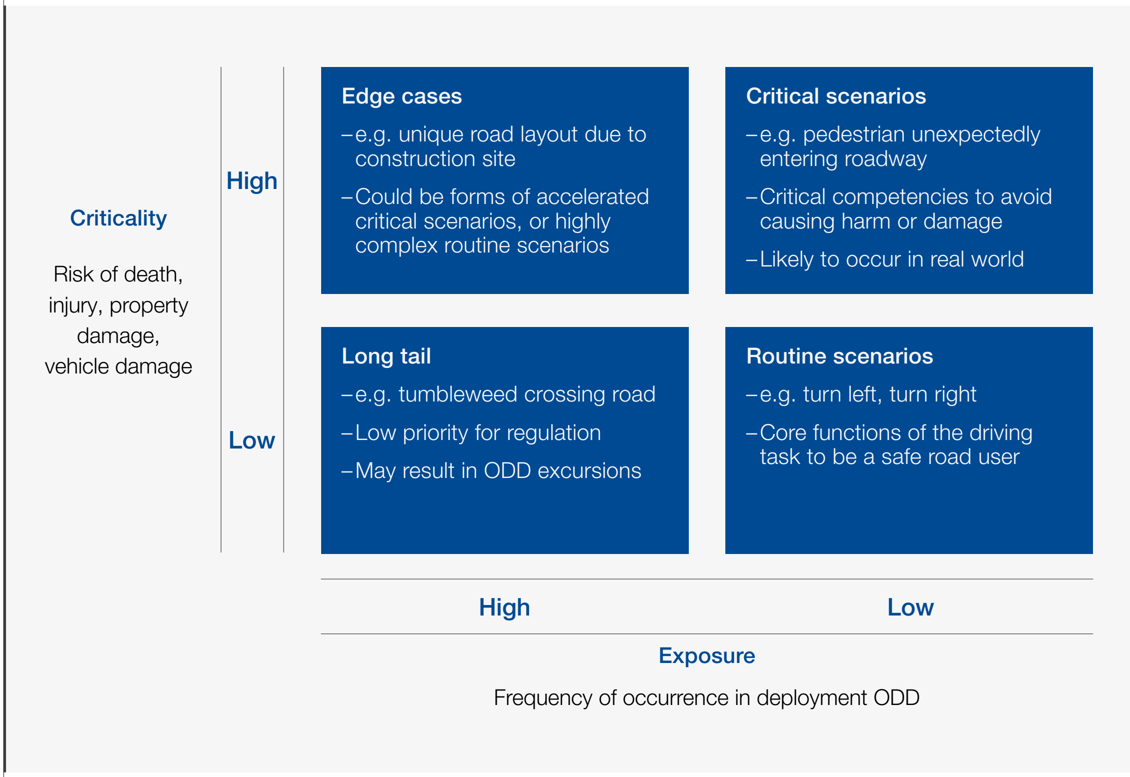

Source:
World Economic Forum / McKinsey & Co. Analysis

An AV system will probably experience all specified qualitative scenarios throughout the safety assurance process, but it is important to prioritize specific parameters (e.g. number of pedestrians, speed of traffic, presence of other vehicles) to test at each stage in the assurance process, to demonstrate the safety and maturity of the system.

For instance, the first milestone's assessment (e.g. see Figure 4) should largely comprise critical scenarios to demonstrate basic competence in avoiding crashes, such as stopping to avoid a pedestrian crossing the road, while edge cases with high degrees of complexity such as unique and rare intersection designs should be tested in later milestones as the system increases in maturity.

Define metrics and pass/fail criteria

With the scenario set established and a corresponding library dataset, the CoE should select measures for each scenario that reflect a successful outcome. These measures will vary across scenarios, and will depend on factors such as the probability of exposure, potential severity and, in some cases, subjective assessments of a successful outcome. Each scenario may have different metrics and outcomes, so it is important to be flexible on this step and work within the structure AV developers have been using to evaluate performance throughout the process.

Sufficient exposure to and successful behaviour in scenarios should be combined with other metrics to form the criteria used to determine when to advance AV systems to the next milestone and eventually to full-scale deployment. To measure scenario exposure, AVs should be aware of each parameter of the ODD in real time, through a combination of on-board sensors and incoming data streams. Therefore, AVs can eventually categorize test data into the high-level behavioural competencies defined by regulators.

In measuring the ODD, AVs can assess ODD excursions, or situations in which the AV leaves its designed ODD. In these situations, the AV can continue operating if safe, execute a minimal risk manoeuvre or return control to a human operator. This is just one potential measure of successful performance and can be combined with additional metrics such as Collison Avoidance Capability, low crash rates and acceptable rules of the road violations.[5]

Structure testing programme

Having chosen the scenarios, testing type and specified metrics for success, the CoE must structure an assessment programme. The approach will vary according to the CoE's preference for in-sourcing this process, as well as the AV developer's established testing programme before entering the regulatory review.

The AV provider and regulator should work together to determine the appropriate assessment modality for each scenario. Additionally, the CoE should create a reporting programme, to enable the AV developers to submit evidence of their performance in simulation, and an audit process if necessary. All modalities will probably be required to assess each AV platform, and considerations for each test type are summarized below:

– **Simulation-based assessment:** Simulations can explore a large number of scenario parameters and values, identifying situations in which the AV platform may fail. The CoE will have to choose between specifying a selection of approved simulation environments and allowing AV developers to demonstrate performance in their existing simulation tools. In any case, the simulation will need to be suitably representative of real-world driving behaviour.

– **Controlled environment testing:** A controlled environment or test track is most useful for structured testing in a repeatable fashion, especially of high-risk, high-frequency scenarios such as pedestrian crossings. AV operators will continue to use structured testing to recreate scenarios throughout their development cycle to validate a range of functions from emergency stops to verifying the fidelity of their simulation for certain scenarios. In this programme, the CoE will have to choose between creating its own test environment, allowing for testing at a third-party site or asking the AV operators to self-certify their own track-based testing.

– **On-road testing:** As with other assessment modes, the CoE can choose either to take a highly structured approach – whereby it creates a formal assessment on the road akin to a driving test – or allow AV operators to self-report their capabilities. This assessment should be representative of the planned driving environment, without asking the AV to follow a single fixed route. Naturalistic on-road testing will probably be a core element of any safety assessment.

Interventions and minimum risk conditions

Despite the best efforts of the CoE to quantify and categorize scenarios in each ODD, the AV systems will likely still encounter unexpected scenarios once deployed. To reduce the chance of unknown scenarios and to validate the AV's abilities to handle such situations, the CoE should test whether the AV can safely reach a minimal risk condition. Hence, the CoE should also consider the need to verify one or more independent mechanisms that successfully execute minimal risk manoeuvres or transition control to human backup operators in unknown scenarios, emergency situations and sustained ODD excursions. Regulators should not necessarily define exactly what these backup mechanisms are at each stage in the process (e.g. two safety drivers at the beginning, one at the end) in order to promote innovative solutions to the problem such as remote operators.

The CoE should require the AV developer to provide a list of potential planning failures and possible mitigations. An AV may have one or more minimal risk manoeuvres[6] that bring the vehicle to a safe minimal risk condition, which may involve an interim degraded operation. The CoE should validate, via design and/or demonstration in a controlled environment, the various minimal risk manoeuvres. These may include manoeuvres such as *transition demand* (request takeover by vehicle operator, whether in-vehicle or remote), *limit function state* (transition to limited operation), *comfort stop* (comfortable transition to end of operation), *emergency stop* (in case of severe or rare system failures) and more.

4.4 Execute – ongoing monitoring

After establishing the testing and approval process, the CoE should continue to oversee the tests as AV developers reach each milestone. AV developers will likely conduct a combination of on-road, controlled test-track and simulation tests, and the CoE should review the results reported by the AV companies and track progress towards each milestone. The CoE should also consider the relevance of the testing programme as an evolving factor and continue to consider the need to modify or expand the scope of the testing programme. For example, if an AV experiences a near miss, or is involved in a collision with another vehicle, the CoE should consider how to capture that data to ensure other vehicles facing the same scenario do not end up in a collision.

The amount and type of data that regulators collect during testing can vary and should be limited to the quantity needed to establish safety while minimizing the risk of disclosing individual AV companies' intellectual property. The data collected should continuously improve the process and add additional scenarios to the library as they arise in naturalistic testing. In addition, the process should be designed to promote continuous testing throughout, without long delays that can hamper progress. Where possible, this reporting process could be automated to reduce friction and resource requirements.

As AVs are deployed in their operating environment, it is important that the CoE monitors the safety performance indicators (SPIs) on an ongoing basis over the life cycle of the AVs. Trust in the technology is earned with good performance over time, and continuous monitoring and feedback loops from the field are one way of assuring ongoing safety. To that end, the CoE should both encourage and require AV developers to report SPIs that demonstrate the effectiveness of their safety case in the field. Incident reports and losses are merely lagging metrics, after-the-event measurements useful for reporting progress but not helpful when judging current or future performance. The CoE should require AV developers to report SPIs that violate the AV's safety claims, such as field failure rates of sensors and components, field failure and error rates of important subsystems such as perception modules, violation of safe clearance limits, concurrent multisensor detection failures and more. For a fuller list of SPIs, refer to a standard such as ANSI/UL 4600.[7]

By requiring the reporting of SPIs on an ongoing basis, and ensuring that the AV developer can update its functionality, the CoE can enable the assurance of safety during the AV's entire life cycle, and continue to develop relevant assessments that are representative of the risks of the driving environment.

Conclusion

After completing this assessment, the regulator should have a clear understanding as to which AV developers are ready to operate commercially in the deployment ODD. Ideally, the AV developer should demonstrate the capability to operate without a safety driver in the vehicle, depending on other backup mechanisms such as minimal risk manoeuvres and/or remote operators to take over control in the event of the AV encountering a rare situation it was not designed to handle.

This framework is intended to provide high-level guidance to regulators desiring to implement an operational safety assessment within their jurisdiction. No one-size-fits-all solution will be sufficient to convince regulators, and the general public, of AVs' safe operation in every city, state and country around the world. Each step in this process should be customized for each locale. However, this customization should be balanced with as much standardization and harmonization is possible at the highest levels of government, as this will better enable AV technology to deploy at scale. To this end, further technical guidance detailing each of the above steps will be available in a forthcoming publication.

The Forum hopes this framework will enable the safe development and deployment of automated vehicles, so that their benefits in delivering a safe, clean and inclusive future of mobility are realized.

Contributors

The World Economic Forum's Safe Drive Initiative is a global, multistakeholder and cross-disciplinary initiative intended to help shape the development of successful autonomous vehicle policy and improve the safety of AV pilots. The project has engaged leaders from private companies, governments, civil society organizations and academia to understand AV policy, identify challenges and define principles to guide future policy solutions. The opinions expressed herein may not correspond with the opinions of all members and organizations involved in the project.

Lead Author

Tim Dawkins
Project Lead, Automotive and Autonomous Mobility, World Economic Forum

We would like to thank our Safe Drive Initiative project community as well as the following contributors for their insights and contributions to this framework:

Atul Acharya
Director, Autonomous Vehicles Strategy, A3 Ventures

Michelle Avary
Head of Automotive and Autonomous Mobility, World Economic Forum

Bryn Balcombe
Founder, The Autonomous Driving Alliance

Xantha Bruso
Manager, Autonomous Vehicle Policy, AAA

Nicola Croce
Technical Program Manager, Safety Pool, Deepen.ai

Thomas Form
Senior Director Autonomous Driving, Volkswagen

Kiran Jesudasan
Transportation Systems Specialist, dRISK

Martin Kellner
Engagement Manager, McKinsey & Company

Siddartha Khastgir
Head of Verification & Validation – Intelligent Vehicles, WMG, University of Warwick

Stefano Marzani
Principal Specialist, Autonomous Driving, Amazon

Dylan Moore
Associate, McKinsey & Company

Mohammad Musa
Chief Executive Officer and Co-Founder, Deepen.ai

Stefan Seltz-Axmacher
Chief Executive Officer and Co-Founder, Starsky Robotics

Sanjit Seshia
Professor, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

Emily Shao
Associate Partner, McKinsey & Company

Barnaby Simkin
Autonomous Driving Standards Manager, NVIDIA

Jann-Eve Stavesand
Head of Consulting, dSPACE

Chess Stetson
Chief Executive Officer, dRISK

Michael Watson
Senior Controls Engineer, NVIDIA

Ismail Zohdy
Program Manager for Self-Driving Transport Strategy, Dubai Road and Transport Authority

Endnotes

1. About CETRAN: https://cetran.sg/ (link as of 29/9/20).
2. NHTSA: A Framework for Automated Driving System Testable Cases and Scenarios: https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/13882-automateddrivingsystems_092618_v1a_tag.pdf (link as of 29/9/20).
3. PAS 1883: Hierarchical Taxonomy for Specifying an Operational Design Domain (ODD): https://www.bsigroup.com/en-GB/CAV/pas-1883/ (link as of 29/9/20).
4. Safety Pool – Scenario Library: www.safetypool.ai (link as of 29/9/20).
5. Driving Safety Performance Assessment Metrics for ADS-Equipped Vehicles: https://www.sae.org/publications/technical-papers/content/2020-01-1206/ (link as of 29/9/20).
6. Safety First for Automated Driving: https://newsroom.intel.com/wp-content/uploads/sites/11/2019/07/Intel-Safety-First-for-Automated-Driving.pdf (link as of 29/9/20).
7. ANSI/UL4600 – Standard for Safety for the Evaluation of Autonomous Vehicles and Other Products: https://ul.org/UL4600 (link as of 29/9/20).

WORLD ECONOMIC FORUM

COMMITTED TO IMPROVING THE STATE OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland

Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744
contact@weforum.org
www.weforum.org