Safety Case Framework: The Guidance Edition

For Reviewers

2021

ZENZIC

SELF-DRIVING REVOLUTION

CREATING A SEAMLESS CAM TESTBED UK

Disclaimer

This report has been produced by HORIBA MIRA, TRL Limited (TRL) and WMG under a contract with Zenzic-UK Ltd (Zenzic). Any views expressed in this report are not necessarily those of Zenzic.

The information contained herein is the property of these organisations and does not necessarily reflect the views or policies of the customer for whom this report was prepared. Whilst every effort has been made to ensure that the matter presented in this report is relevant, accurate and up-to-date, Zenzic and/or any of the authors of this report cannot accept any liability for any error or omission, or reliance on part or all of the content in case of incidents that may arise during trialling and testing. In addition, Zenzic and/or any of the authors of this report cannot accept any liability for any error or omission, or reliance on part or all of the content in another context.

When in hard copy, this publication is printed on paper that is FSC (Forest Stewardship Council) and TCF (Totally Chlorine Free) registered.

For further information on this report please contact the Zenzic team

info@zenzic.io
zenzic.io

Table of acronyms:

ADS	Automated Driving System
ALARP	As low as reasonably practicable
ATA	Attack Tree Analysis
AV	Automated Vehicle
CAM	Connected and Automated Mobility
CAV	Connected and Autonomous Vehicle
CCA	Cause-Consequence Analysis
CCAV	Centre for Connected and Autonomous Vehicles
DDT	Dynamic Driving Task
FMEA	Failure Modes and Effects Analysis
FTA	Fault Tree Analysis
GSN	Goal Structuring Notation
HAZOP	Hazard and Operability Study
MRC	Minimal Risk Condition
MRM	Minimal Risk Manoeuvre
ODD	Operational Design Domain
SOTIF	Safety of the Intended Function
STPA	System Theoretic Process Analysis
STRIDE	Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege
TVRA	Threat, Vulnerability and Risk Analysis
UNECE	United Nations Economic Commission for Europe

About CAM Testbed UK

CAM Testbed UK, coordinated by Zenzic, is a collection of world-class testing and development facilities for connected and self-driving technologies. Supported by the UK government's Centre for Connected and Autonomous Vehicles (CCAV), it is the hub for excellence in testing and assurance for CAM and is built on a long history of testing and assurance of vehicles in both private and public scenarios.

About the authors

Richard Hillman – HORIBA MIRA:

Richard is a Principal Engineer at HORIBA MIRA, working within the connected and autonomous vehicles (CAV) department. He is a specialist in systems engineering and safety engineering, with experience of developing advanced driver assistance systems (ADAS) and of autonomous vehicle trials. This has included being responsible for safety management on high profile projects such as HumanDrive, together with research on test methodologies and simulation.

Paul Wooderson – HORIBA MIRA:

Paul is Cyber security Chief Engineer at HORIBA MIRA. He is a Chartered Engineer with 20 years' experience in security engineering in the automotive and semiconductor industries. Paul is a UK expert to the international ISO and UNECE working groups, developing new standards and regulations for cyber security engineering and software updates for road vehicles. He has played a lead role in the recent cyber security related collaborative research projects ResiCAV, 5StarS and UK CITE.

Dr Siddartha Khastgir – WMG:

Siddartha is the Head of Verification & Validation of CAV at WMG, University of Warwick, leading various collaborative R&D projects with industrial and academic partners. His research areas include test scenarios generation, system safety, simulation based testing and safe AI, among many others. He represents the UK on various ISO Technical Committees and is the lead author for new ISO standards for Low-Speed Automated Driving (LSAD) systems (ISO 22737) and Taxonomy for ODD (ISO 34503). He is also the project lead for ASAM standardisation project - OpenODD. He has received numerous national and international awards for his research contributions, including the prestigious UKRI Future Leaders Fellowship - a seven year grant focused on safety evaluation of CAVs.

Chris Fordham – TRL:

Chris is a Principal Safety & Risk Consultant at TRL and has over 5 years' of safety and risk management experience in the automated vehicle and road safety fields. Chris has authored and independently reviewed a number of safety cases, operational risk assessments and operational guidance documents for automated vehicle trials to ensure acceptable operational safety. His automated vehicle safety expertise has also fed into industry standards and guidance, including BSI PAS 1881 and the Zenzic Safety Case Framework Report V2.0.

Oliver Howes – TRL:

Oliver has over five years professional engineering experience within automotive engineering consultancy, vehicle and ICE testing, manufacturing and now transportation. Oliver has developed the Smart Mobility Living Lab's (SMLL) customer onboarding process, including a method of assessing a customer safety case, and led the development of a systems safety case framework for the ServCity project. Oliver also contributed to the Zenzic Safety Case Framework Report V2.0, and has been the technical lead for multiple real world trials within the SMLL facility in London.

Dr Anneka Lawson – TRL:

Anneka completed her doctoral studies in Transport Engineering related to cycling safety and network modelling at the University of Dublin before a move to London, where she spent a number of years working in transport policy research related to road safety and other transport issues at the RAC Foundation. She now works as a safety and risk consultant, specialising in CAV technologies, working with clients such as the Centre for Connected and Autonomous Vehicles (CCAV) and Zenzic, and on Innovate UK funded projects such as ServCity and Project Endeavour.

Will Perren – TRL:

Will is a Safety and Risk Consultant at TRL with expertise in CAV trialling, standards and regulations. He has authored a number of operational safety cases and risk assessments to ensure CAV trials are managed safely. Will also works closely with CAV stakeholders including testbeds, local and road authorities to develop harmonised approaches to manage CAV trials safely. His operational safety expertise has fed into the industry standards and guidance, including the Zenzic Safety Case Framework Report V2.0, where he was the lead author.

Acknowledgements

Zenzic Project Lead: Bhavin Makwana
Project Manager: Mili Naik

Special thanks to the Advisory Group for guiding and shaping this work. The Advisory Group comprised: Centre for Connected and Autonomous Vehicles (CCAV), Driver and Vehicle Standards Agency (DVSA), Vehicle Certification Agency (VCA), Oxbotica, Arrival, ZF, Oxford County Council, TfL, BSI and Highways England.

Special thanks also to Camilla Fowler (TRL) for her input and guidance throughout the project.

Contents

Key

Throughout the document, examples have been highlighted in grey, while key evidence or assurance that Safety Case Reviewers are advised to look for have been highlighted in blue. The latter is to make it easier to readers to identify the most important aspects and prioritise their focus accordingly, and hence particular emphasis should be placed on ensuring that trials are compliant with these blue clauses.

EXAMPLES AND TEMPLATES

KEY REVIEWER POINT

Key evidence or assurance that Safety Case Reviewers are advised to look for have been highlighted in blue.

The guidance applies to all trials (typically these are where there is limited or no control of the surrounding environment but a safety driver is able to override via traditional controls) unless tagged as outlined below. Throughout the document, sections have been tagged to show the types of trials this section of guidance is applicable to.

PRELIMINARY TRIALS

Safety cases for relatively simple trials within a highly-controlled environment and with a safety driver who is able to override via traditional controls would be deemed as preliminary trials. In general, these trials would only have these tagged requirements. This would also apply to trials involving manual driving e.g. for data collection. Guidance tagged as 'preliminary trials' is also applicable to all other types of trials.

ADVANCED TRIALS

Safety cases for trials without a conventional safety driver (e.g. trials that use a remote safety operator, have limited controls or are without a safety operator at all) would be deemed an advanced trial. This is additional to all other guidance provided, including the sections marked 'preliminary trials'.

Foreword

Safety is a critical enabler to realising the wider social and economic benefits of connected and automated mobility (CAM). With the market forecast to be worth in excess of £650bn globally by 2035 there is a critical need to ensure that consumers and businesses have confidence in the development and delivery of future transportation systems. Without this it will be impossible to unlock the productivity, efficiency and inclusion benefits of CAM.

The Safety Case Framework: The Guidance Edition reports developed by CAM Testbed UK will enable safety to be delivered in a high quality and consistent manner across testing and trial deployments of CAM services in the UK. By bringing together expert knowledge from the organisations who have been leading on development and delivery of connected and automated technologies over the last five years this guidance provides a concise and authoritative guide to best practice in this emerging market.

Furthermore, drawing on guidance from CCAV's *Code of Practice* and BSI's *CAV Standards Programme*, this safety case guidance is a first in terms of collating and summarising all the key learning of standards and best practice documentation into a single easy to access document. While we do not aim to replace the in-depth knowledge contained in other documents, we do aim to make it easier for both those developing safety cases and those reviewing them to understand how and where to apply best practice.

Longer term there are still significant questions to answer in terms of the route from early trials to full scale deployments with significant changes needed in certification processes and legislation. But for now, we hope to create a more effective dialogue between organisations who need to test and those who want to support CAM testing so that government and industry can rapidly iterate on best practice and accelerate the self-driving revolution.

Mark Cracknell
Head of Technology
Zenzic

PRELIMINARY TRIALS

1.0 Introduction

As connected and automated mobility technology develops, there is an increased demand to test the performance and capabilities of such technology in a range of different operating conditions and environments.

The primary aim of this safety case guidance is to support organisations tasked with reviewing safety cases for trials at CAM Testbed UK, in order to ensure good safety practice is applied to all trials and in order to align the expectations of testbeds.

However, it is important to note that this guidance does not seek to impose any particular methodology or format, as it must be acknowledged that there is no single approach universally recognised to be demonstrably superior to others. It is also important to bear in mind that most companies will already have corporate policies on risk management, and a trial that has tested elsewhere previously would already have its own safety case format. This openness is encouraged and testbeds should remain flexible in the safety cases they accept, to avoid trials with a pre-existing format having to commit time and resources to modify it to fit an arbitrary template. Efforts have been made to accommodate the known requirements of testbeds within CAM Testbed UK within this guidance.

There are a range of documents, such as regulations and standards, that are already applicable to trials taking place within the UK, and these are referenced as appropriate. However, this guidance document does not seek to impose further requirements over and above those already imposed.

Instead, the aim is to help Safety Case Reviewers by providing explanations and examples to help illustrate possible best-practice solutions, and to link the areas covered by the separate regulations and standards together such that the overall picture of how they contribute to the safety case becomes apparent (see Figure 1.1). Trialling organisations are welcome to make use of the examples provided wherever they add value, but they are not mandatory to follow. Examples have been highlighted in grey to make them identifiable throughout.

This guidance has been created based upon extensive research into existing methodologies, including three industry workshops attended by stakeholders with a wide range of perspectives. It also draws upon the significant experience that the team of technical authors have relating to trial safety management. The style of the document has been informed by stakeholder feedback, including regular steering meetings with an advisory group of industry experts, in order to arrive at something that is proportionate to the level of complexity and risk and that is of assistance to both Safety Case Creators and Reviewers.

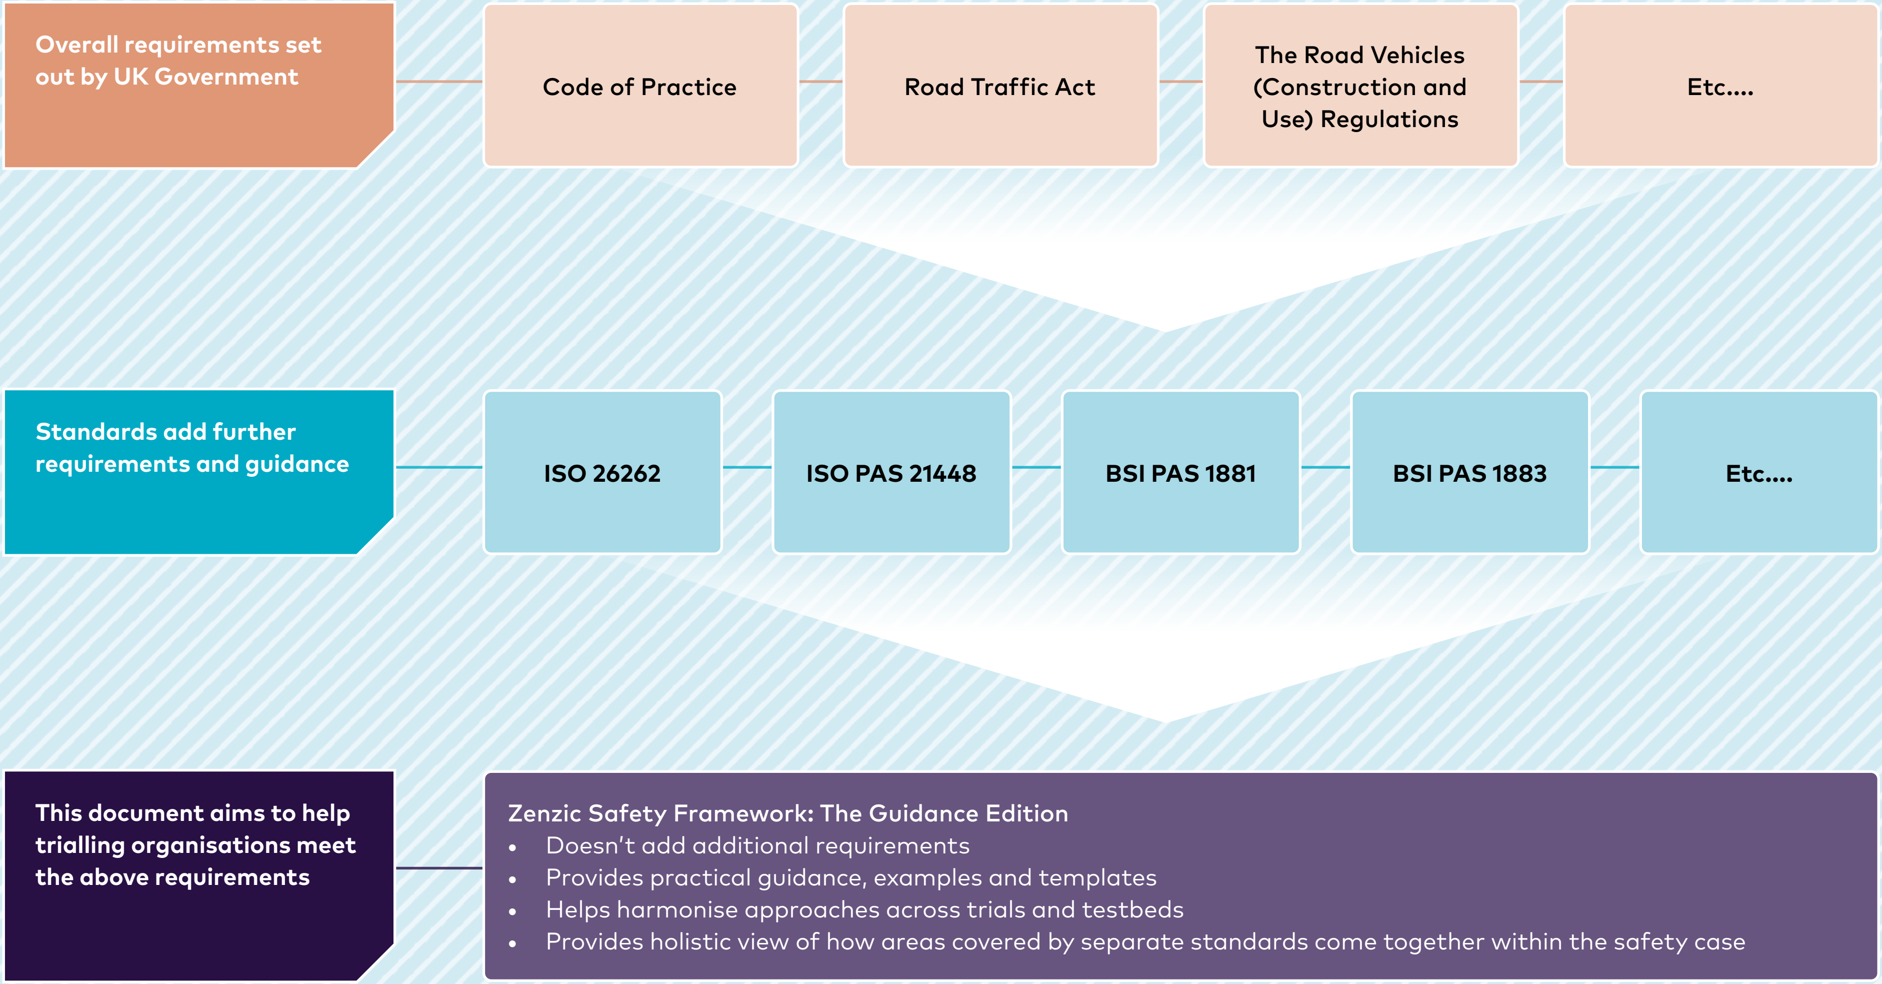

Overall requirements set out by UK Government

Code of Practice

Road Traffic Act

The Road Vehicles (Construction and Use) Regulations

Etc....

Standards add further requirements and guidance

ISO 26262

ISO PAS 21448

BSI PAS 1881

BSI PAS 1883

Etc....

This document aims to help trialling organisations meet the above requirements

Zenzic Safety Framework: The Guidance Edition

- Doesn't add additional requirements
- Provides practical guidance, examples and templates
- Helps harmonise approaches across trials and testbeds
- Provides holistic view of how areas covered by separate standards come together within the safety case

Figure 1.1: Illustration of where this document sits within the wider landscape of regulations and standards

The way this document is applied will depend upon the inherent complexity and risk involved in a trial. For example, it would not be proportionate for a trial on a proving ground with a safety driver present in the vehicle to go through the same safety processes as a trial involving a remotely supervised vehicle operating upon a public road. Section 2 examines how trialling organisations should approach determining what level of analysis is appropriate, and this document further aims to aid proportionality by clearly denoting which sections would be applicable to 'preliminary trials' (trials in a highly-controlled environment with a high level control provided by a safety driver in the vehicle, and trials featuring manual control of the vehicle) and 'advanced trials' (with little or no control of the environment and without a safety driver able to make manual interventions using conventional driver controls). It should be noted that many trials will progress over time (e.g. a trial may start off upon a proving ground before progressing to public roads, or start with a safety driver in the vehicle before progressing to a remote safety operator). As such, it is important to remember that sections of the guidance that were not applicable in the early stages of a trial may become applicable in the later stages.

PRELIMINARY TRIALS

Safety cases for relatively simple trials within a highly-controlled environment and with safety driver who is able to override via traditional controls would be deemed as preliminary trials. In general, these trials would only have these tagged requirements. This would also apply to trials involving manual driving e.g. for data collection. Guidance tagged for preliminary trials is also applicable to all other types of trials. The key applicable sections would be the risk assessment (Section 4.1) and method statement (Section 4.2.2.1).

ADVANCED TRIALS

Safety cases for trials without a conventional safety driver (e.g. trials that use a remote safety operator, have limited controls or are without a safety operator at all) would be deemed an advanced trial. This is additional to all other guidance provided, including the sections marked 'preliminary trials'. These safety cases would generally require consideration of all sections of this document, including those marked 'advanced trials'.

Although this guidance document is directly focussed upon trials within CAM Testbed UK, the overall principles of best practice remain the same for any test route, and thus it is hoped that this guidance will also provide value to trials elsewhere in the UK, and indeed overseas. The scope is, however, limited to trials for the purpose of research and development, and should therefore not be taken as a solution for the very different challenge of providing safety assurance for full deployment of a commercially available CAM solution.

This provides differentiation from other initiatives such as CertiCAV and CAV PASS, which are targeted at developing an approval process to allow commercial deployment and therefore place an emphasis on system safety assurance rather than on operational safety measures such as the use of a safety driver. However, it should be noted that there is some overlap regarding 'advanced trials' of automated vehicles on public roads without a safety driver, as CAV PASS will initially focus upon such trials as an intermediate step towards the ultimate goal of providing safety assurance for full deployment (CCAV, 2019a).

PRELIMINARY TRIALS

1.1 Terms and definitions

This document makes use of Version 3.0 of the CAV (Connected and Autonomous Vehicle) Vocabulary published by BSI (2020). Readers should therefore refer to this where further definitions are required. In general, however, this document attempts to introduce terms and acronyms such that cross-referencing is not required; this is particularly so for any terms not contained within the BSI CAV Vocabulary. Readers should also be aware of SAE J3016 (SAE, 2018a), which provides a taxonomy of definitions that this document maintains consistency with.

It should be noted in particular, however, that this report uses the term '**safety driver**' to refer specifically to a person who is physically present in the vehicle, able to observe the surroundings in the manner of a regular driver, and able to assume control of the vehicle via a conventional set of driver controls. '**Safety operator**', on the other hand, is a broader term that encompasses safety drivers but also other solutions such as remote safety operators or safety operators who are in the vehicle but only have access to limited controls such as an emergency stop button.

2.0 Principles and background

This section examines what components a safety case would typically have and how to ensure the level of detail is proportionate to the complexity and risk presented by the nature of the trial.

2.1 Safety case types and structure

PRELIMINARY TRIALS

2.1.1 The operational safety case, system safety case and security case

The Code of Practice for Automated Vehicle Trialling requires that all trialling organisations should develop a detailed safety case for any trialling they wish to undertake in the UK (CCAV, 2019b).

A safety case is an essential tool to demonstrate how safety and security has been assessed and managed, and can be categorised into three interdependent areas: **system safety, operational safety** and **security**.

A **system safety** case focuses on the safety of the system under test, including 'functional safety' (i.e. managing risks resulting from potential system faults) and 'safety of the intended function' (i.e. managing risks due to inherent design limitations that are present even when the system is functioning as intended). The purpose of the system safety case is to document system safety assessments and demonstrate that the system presents a level of safety that is proportionate to the testing proposed. The systems safety case is intrinsically linked to the system and can be used for multiple trials operating in the same ODD (Operational Design Domain, as defined in Section 3).

It is worth noting that a 'system' can be defined at various levels; for example, when developing an integrated transport network, the entire infrastructure could be referred to as a system, with the vehicles within that system being subsystems, whereas a developer of a driver assistance feature such as traffic sign recognition may regard this feature as the system, with the vehicle being a supersystem. Within this document, the system is taken to be the interconnected collection of physical, electromechanical, electronic and data elements of the vehicle, including automated driving system (ADS) that enables full or partial automation. In addition to hardware and software onboard the vehicle, this would also encompass any offboard subsystems that directly facilitate automated driving.

An **operational safety** case is a structured body of evidence that considers the interaction of the test vehicle(s) with the operating environment, including the route, safety driver or operator, passengers and other road users. The main purpose of an operational safety case is to demonstrate that the vehicle can operate safely within the defined environment and to provide appropriate evidence and mitigations proportionate to the level of risk posed. The operational safety case is location- and time-specific, and therefore should be bespoke for each trial.

For trials of early prototypes, it will frequently be the case that trial safety is primarily dependent upon operational safety measures such as the ability of a safety driver to intervene, rather than being primarily dependent upon system safety assurance. This is because sufficient evidence that the system performance is acceptably safe, such that operational safety measures are unnecessary, will not be available until a significant volume of development and testing has been undertaken.

A **security** case provides evidence that risks presented by harmful actors accessing or affecting any of the trial equipment, including the automated driving system (ADS), have been analysed and mitigated. This includes risks presented by physical access or via electronic and telecommunications means (cyber security).

There is a strong link between these safety case elements, and combined, they form the complete safety case required to ensure trials are conducted in a safe manner. The safety case should remain a live document and be updated to account for changes such as previously unknown hazards being uncovered or changes to the scope of the trial, and should demonstrate that risks have been managed such that they are as low as reasonably practicable (ALARP). This provides assurance to stakeholders such as road operators, landowners, insurers and members of the public that the system can operate safely within the vicinity of other parties and infrastructure.

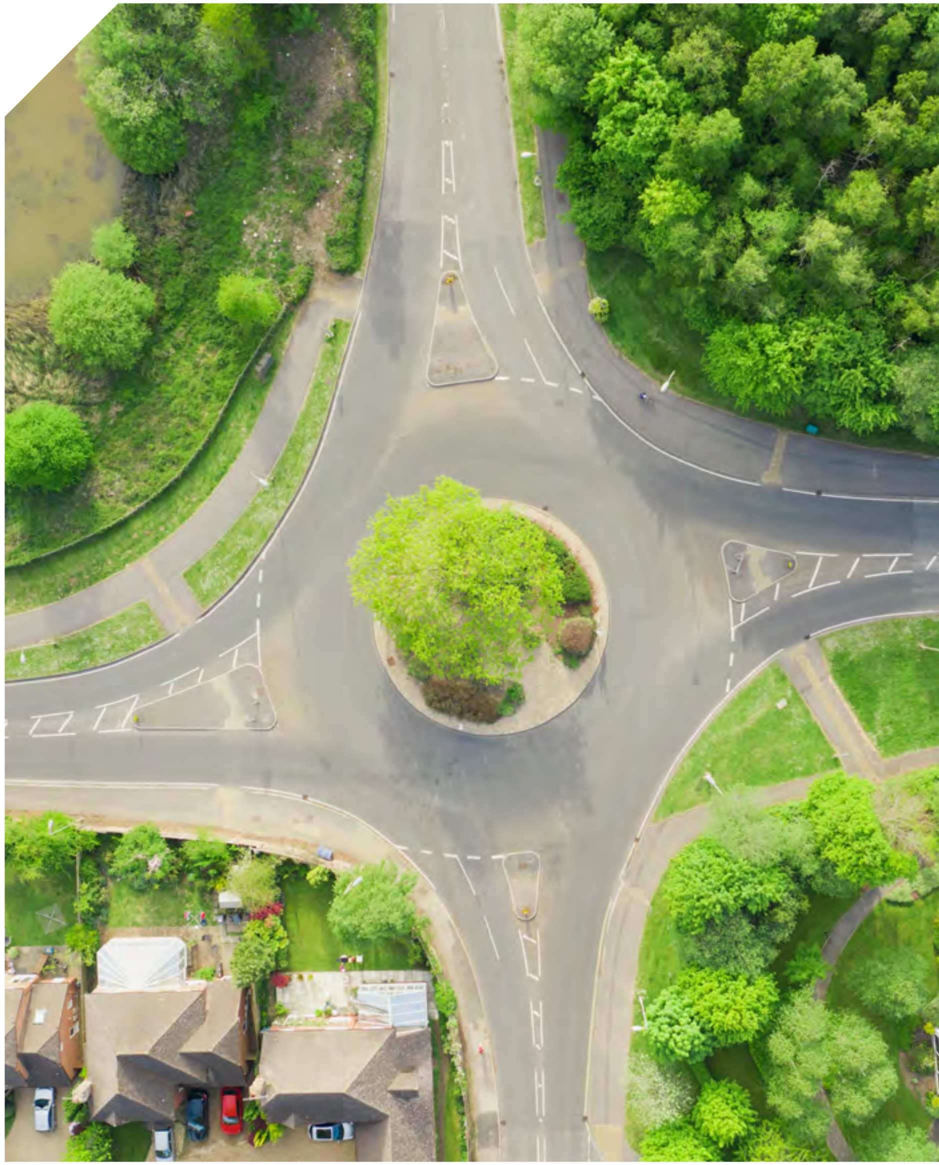

2.1.2 The safety argument

An essential component of the trial safety case will be the **safety argument**; this describes a means of justifying and documenting how all the evidence presented within the safety case, when taken together, supports the overall goal of the trial being acceptably safe. Without a coherent safety argument, the safety case would merely be a mass of information, with no means to understand how it fits together and no means to identify any gaps where there is insufficient evidence to demonstrate the safety of a particular aspect of the trial.

Goal Structuring Notation (GSN) is a widely used approach to presenting a safety argument, allowing it to be displayed graphically with the overall safety goal at the top of the diagram and other sub-goals arranged beneath to support it, with these being progressively broken down until sufficient granularity is reached where specific pieces of evidence can be provided in support (GSN, 2018).

There is no obligation to use GSN, and the safety argument could be conveyed by other means, e.g. by descriptive text. Furthermore, the examples to the right and overleaf should be tailored to suit a particular trial; for example, some trials may have little or no reliance upon system safety, or may have other operational safety measures included beyond those in the example. Nonetheless, it is important for the safety case to include some means to explain how the evidence fits together to form a complete and cohesive safety argument.

EXAMPLES AND TEMPLATES

Figure 2.1 proves an example of what the top-level safety argument might look like, with the overall 'safety goal' (rectangle) of the trial supported by achievement of acceptable operational safety, system safety and security. The three supporting items are all 'modules' as indicated by the small tab on the top left; this means that a further level of safety argument exists beneath this, such modularity aiding readability for complex safety arguments.

An example of how the operational safety case could be further broken down within the next level of decomposition is shown within Figure 2.2, where the top level goal (identical to the module it sits beneath, from Figure 2.1) is supported by multiple other safety goals, with decomposition being continued until the underlying pieces of evidence are reached (circles). The structure of the diagram allows all stakeholders in the safety case to understand the link between the underlying evidence and the overall safety goal of the project being acceptably safe.

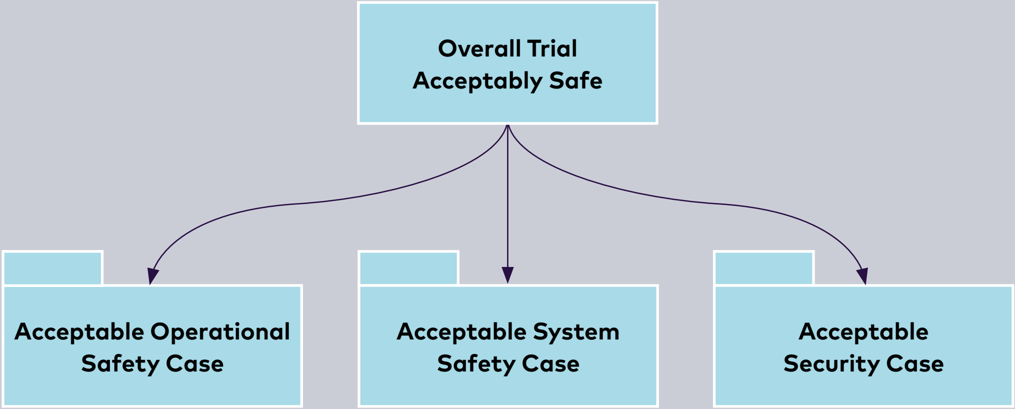
Overall Trial Acceptably Safe
Acceptable Operational Safety Case
Acceptable System Safety Case
Acceptable Security Case

Figure 2.1: An example of a top-level safety argument described using Goal Structuring Notation (GSN)

EXAMPLES AND TEMPLATES

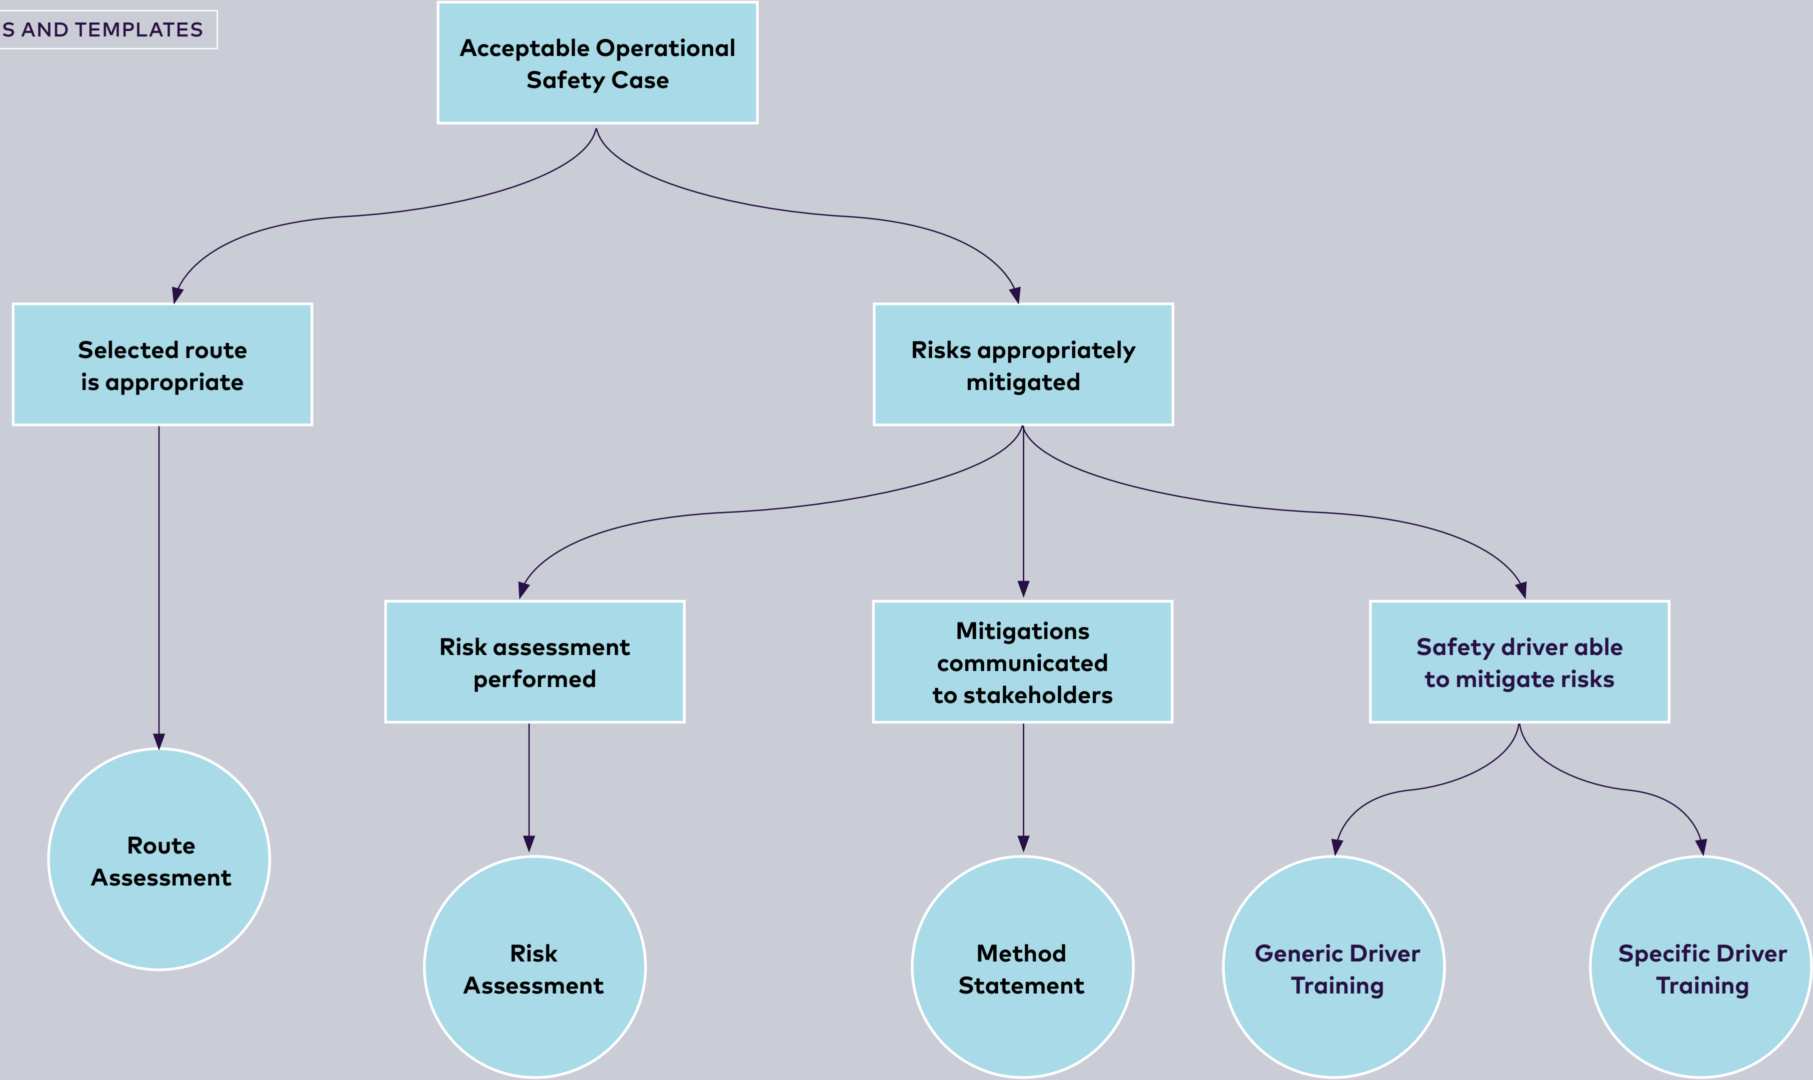
Acceptable Operational Safety Case
Selected route is appropriate
Risks appropriately mitigated
Risk assessment performed
Mitigations communicated to stakeholders
Safety driver able to mitigate risks
Route Assessment
Risk Assessment
Method Statement
Generic Driver Training
Specific Driver Training

Figure 2.2: Example of how evidence could be shown to support the operational safety case, using GSN notation

KEY REVIEWER POINT

The 'safety argument' contained within the safety case documentation should be sufficiently clear for all stakeholders to be able to understand the role of each component document of the safety case in demonstrating the overall safety of the trial. Reviewers should satisfy themselves that the safety argument makes logical sense and contains no clear gaps.

2.1.3 Trial risk factors

The level of risk posed by a trial depends upon three broad, but independent, risk factors:

01 TRIAL ENVIRONMENT

02 SAFETY OPERATOR

03 VEHICLE/ SYSTEM

The level of risk posed depends on the maturity and reliability of the **vehicle** and **automated driving system** (ADS), the ability of the **safety operator** to intervene where necessary and the level of control or predictability of the **trial environment**.

Although production automated vehicles (AVs) ready for commercial deployment would need to have a sufficiently high level of system safety that continuous human oversight is unnecessary, this is not typically the case for research trials of such vehicles, and indeed would not be the case for production systems until their development cycle nears completion. This is due to the large volume of evidence that would be needed to prove they are capable of safe operation, bearing in mind the complexity of the systems, their operating environment and the manoeuvres they would be required to perform.

This guidance document therefore assumes that operational safety will be the primary means of controlling risk during the majority of trials on CAM Testbed UK, and will be a significant factor in all trials. However, the close interaction between the level of trust in the system and the demands for operational safety means that the system safety case and the security case are key influences on the nature of the operational safety case.

2.1.4 The safety case structure

The information contained within a safety case will vary significantly depending upon the nature and complexity of a trial. Furthermore, trialling organisations may choose different approaches to subdivide this information into separate documents; some safety cases could even provide all the necessary information within a single document, although this would make updates challenging as changes to one area would require up-versioning of the whole safety case. More typically, the safety case would consist of multiple documents, each covering a particular aspect, one of which would need to describe how the separate documents fit together (the safety argument, as described in 2.1.2, being key to this).

All safety cases should have common goals of demonstrating an appropriate level of due diligence and providing evidence of control for the trial risk factors. BSI PAS 1881 (2020) specifies the requirements of an operational safety case and reflects current good practice across the industry.

The level of detail required within the safety case should be proportionate to the complexity of the trial and level of risk posed. Section 2.3.1 provides further guidance on how to determine trial complexity.

EXAMPLES AND TEMPLATES

Due to the variance in trials and in their safety case needs, it is not possible to specify a single format, but an example of a possible safety case structure would be:

Summary

- **Overview** of the purpose and scope of the trial
- **Summary of the technology**
- **Safety argument** to explain how the safety evidence fits together
- **Trial monitoring, reporting and continuous improvements processes** including incident reporting and change control
- **Emergency response and crisis communication plans**
- **Stakeholder consultation**

Supporting

- **ODD definition**
- **Operational risk assessment(s)** – separate risk assessments may be needed for separate activities
- **Method statement(s)** – there may be separate method statements to cover different phases of the trials
- **Route selection and assessment**
- **System safety case** including safety analysis of the system, and also simulation and physical testing
- **Security case** including physical, cyber and personnel – both risk assessment and communication of control measures
- **Publicly available safety case**

To illustrate a possible structure for how the information could be grouped together, the first six items (boxed out in blue) could be combined within a single safety case summary document. This would act as the main, central component of the safety case, providing key information on the project background and processes and cross referencing other supporting documents, shown in the light grey box. However, the information could be grouped differently, and trialling organisations should tailor the structure of their safety cases according to their needs.

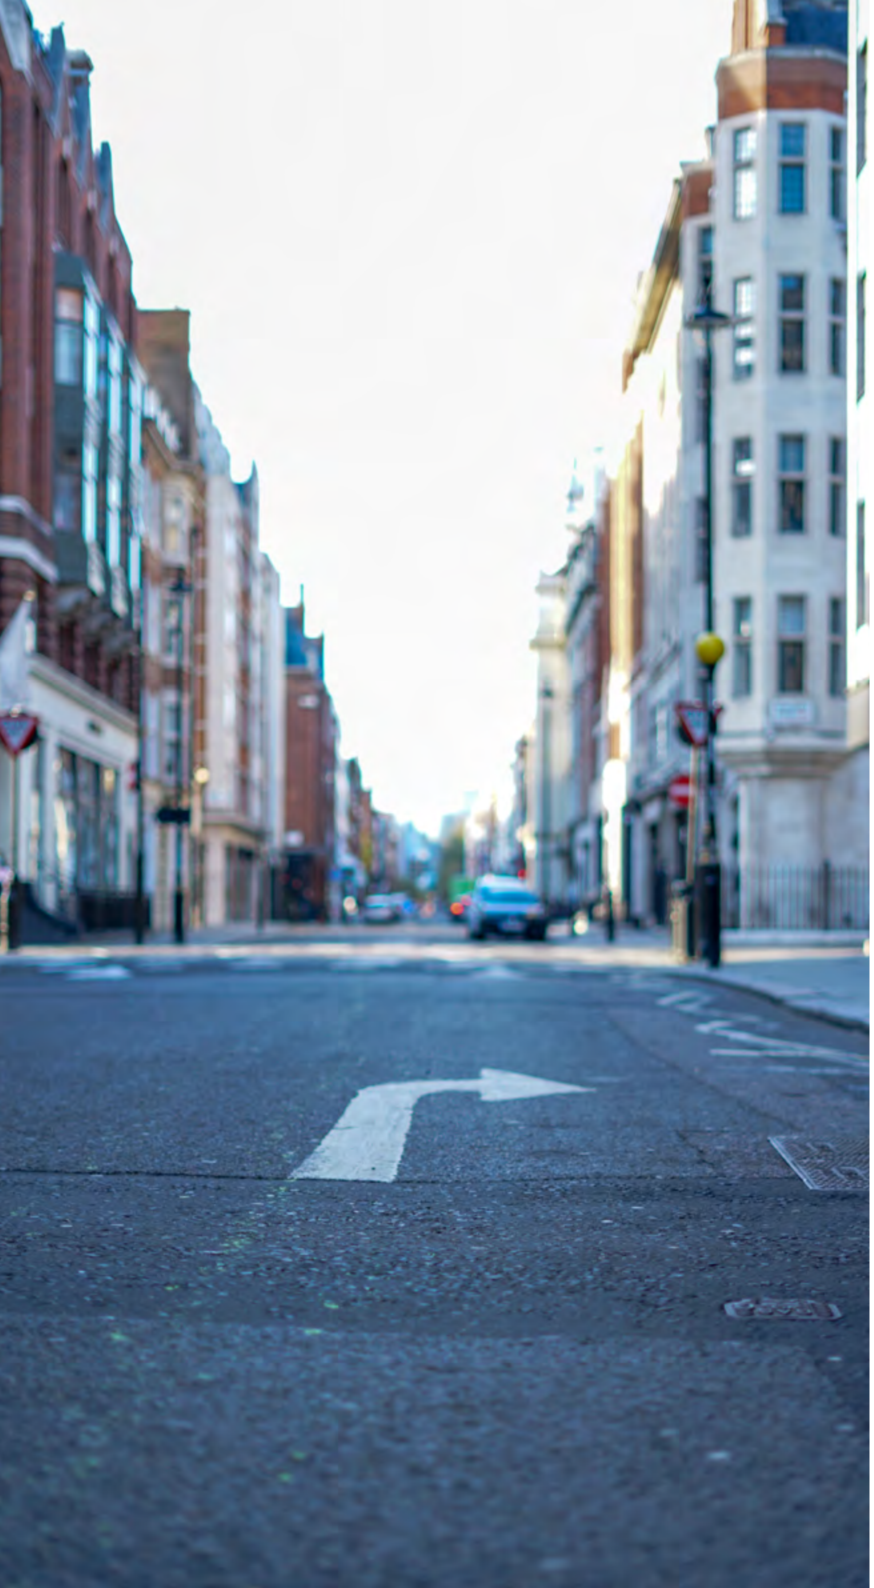

2.2 Who is a Reviewer?

A Safety Case Reviewer can be any stakeholder with an interest in the safety of trial operations. This could include:

- **Test facilities (including proving grounds and public test environments)**
- **Insurers**
- **Highway authorities**
- **Road operators**
- **Landowners**
- **Leaseholders**

Depending on the test location, more than one organisation may need or wish to review the safety case prior to trialling. Within each organisation the safety case may need to pass through multiple departments or subject specialists for review. A reviewer may not be a specific job role within an organisation.

The reviewer is not required to provide a full audit of all safety evidence or assume any legal responsibility through certifying a trial as safe, but is likely to have a desire or obligation to have been shown sufficient evidence to be satisfied that the trialling organisation is applying an appropriate level of diligence in managing the safety of the trial.

The reviewer should, where practicable, be independent from all organisations involved in the trial management and safety case creation. However, it is acknowledged that full independence is not always possible due to collaboration between organisations within past, present or prospective projects or consortia. Where this is the case, reviewing organisations should take steps to ensure impartiality, e.g. by assigning the task to a person or department not connected to the collaboration.

2.2.1 What information does a reviewer review?

As an operational safety case is trial and location-specific and typically forms the main method of safety assurance within AV trials, the review should primarily focus upon the operational safety case. However, as specified in BSI PAS 1881, high-level system information is also required within the operational safety case, as an understanding of the fundamental nature of the technology is necessary if the operational risks are to be understood.

If the safety case seeks to provide evidence that human oversight is not required, or can be provided by a remote operator or other such solution that gives lower control than a safety driver in a vehicle, more detailed evidence of system safety and security would need to be documented.

Depending on the complexity of a trial, the level of detail included in a safety case can vary as the depth of information required changes. The reviewer could be provided with either the same safety case used internally by the trialling organisation or an abridged version which captures the key information needed for reviewers to have a proportionate level of oversight.

2.2.2 The role of a Safety Case Reviewer

It is likely that a Safety Case Reviewer may not be a subject matter expert in the area of technology being tested or in running CAM technology trials. The role of the reviewer is to provide a high-level check as opposed to a full technical review of the safety case, and hence they should ultimately 'accept' that it is appropriate to proceed if they are provided with sufficient evidence of due diligence being applied with regard to safety, but not provide a full 'approval' or 'certification' of the trial safety case. Responsibility for ensuring the safety case is complete and accurate, and that risks are appropriately managed, remains with the trialling organisation.

The review should include a high-level check of included documents to ensure key areas are considered, key risks are addressed, and relevant processes are in place. This may include those listed in Section 2.1.4, and more information is provided on appropriate safety case documentation throughout this guidance.

Many reviewers will be familiar with the trial environment. The reviewer should therefore include within their assessment whether the planned testing activities are appropriate for the test area. The reviewer may be able to provide testbed-specific information to the trialling organisation during the early stages of safety case development, to ensure that the trial is appropriate to the test location and that all local requirements are incorporated.

A reviewer may choose to engage a third party to undertake a safety case review, or to provide consultancy to support a review. For example, although the reviewer does not need to have expertise in the technology being tested to assess an operational safety case, for more advanced trials where the system safety forms a key part of the safety case, more specialist technical knowledge is likely to be required.

EXAMPLES AND TEMPLATES

Figure 2.3 shows a possible high-level process for trialling organisations engaging with a Safety Case Reviewer and the reviewer providing safety case acceptance.

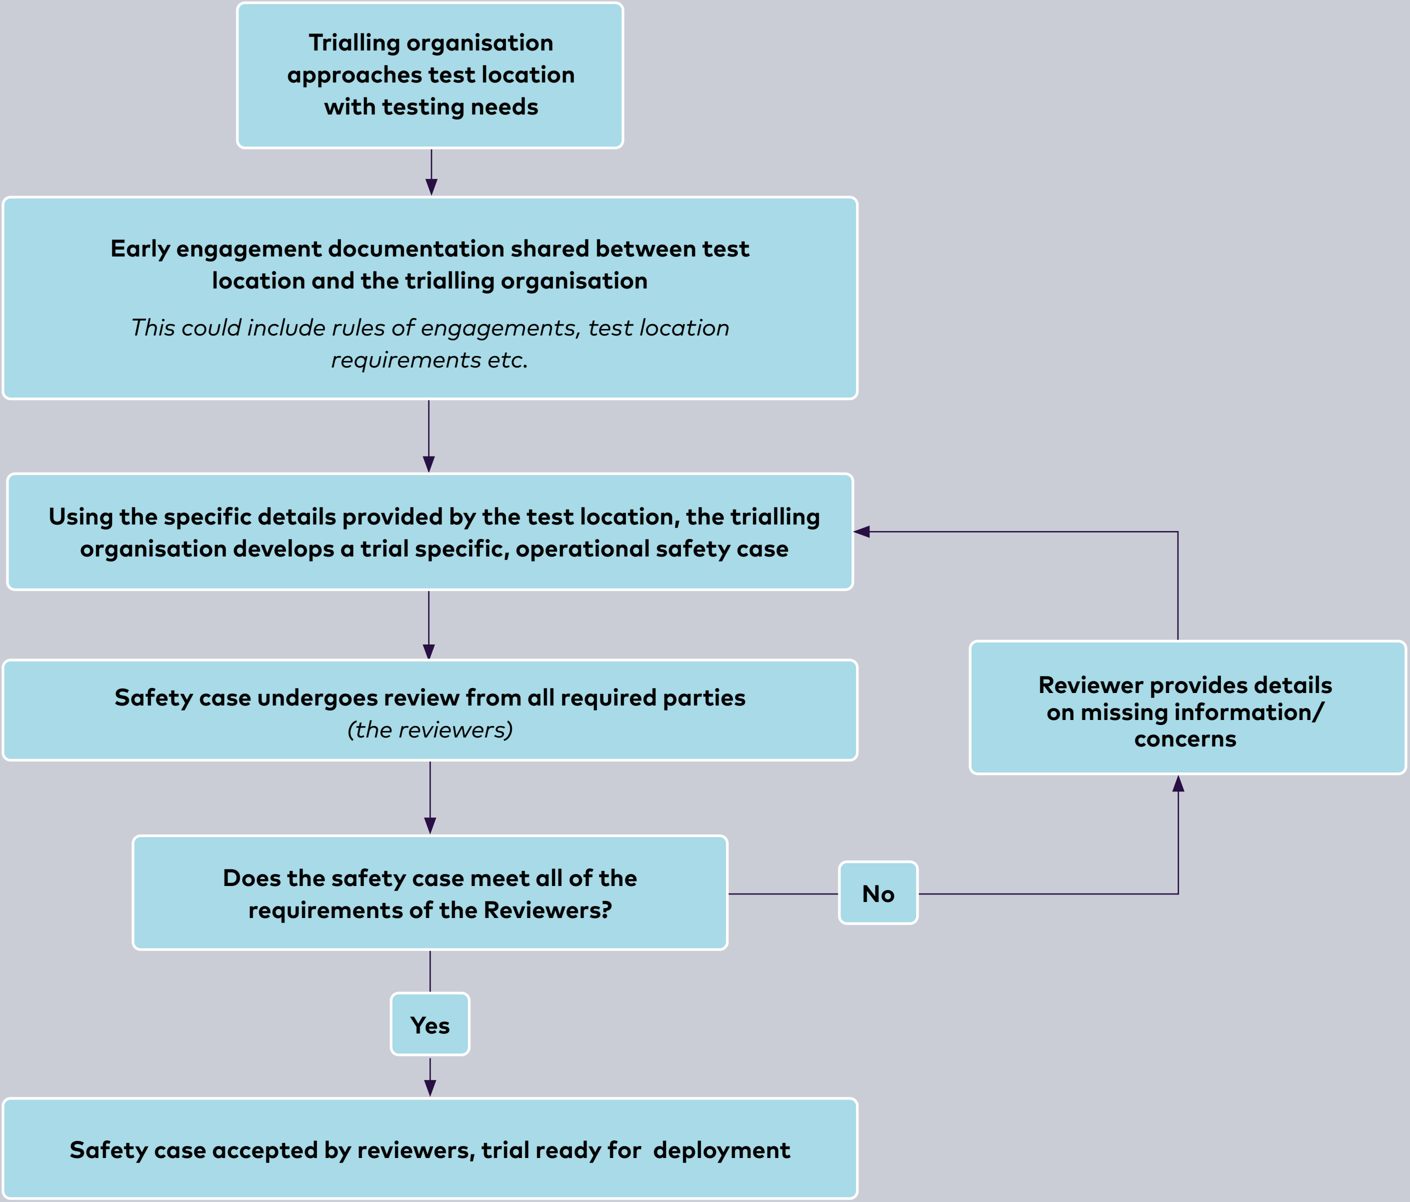
Trialling organisation approaches test location with testing needs
Early engagement documentation shared between test location and the trialling organisation
This could include rules of engagements, test location requirements etc.
Using the specific details provided by the test location, the trialling organisation develops a trial specific, operational safety case
Safety case undergoes review from all required parties (the reviewers)
Does the safety case meet all of the requirements of the Reviewers?
No
Reviewer provides details on missing information/ concerns
Yes
Safety case accepted by reviewers, trial ready for deployment

Figure 2.3: Example process flow for safety case acceptance

2.2.3 How can a Reviewer assist a trialling organisation?

It is recommended that trialling organisations should consult with all necessary reviewers at an early stage in trial planning to ensure smooth and effective safety case development; this will allow appropriate time for all parties to prepare for the testing activities. Safety Case Reviewers can provide early assistance to trialling organisations by informing them of any pre-requisites, policies and local information for the test area that would assist in trial planning and safety case development.

Table 2.1 lists example information that a Safety Case Reviewer could provide to a trialling organisation to assist in the safety case development. This list is not exhaustive and not all entries will be applicable to all Safety Case Reviewers.

Assisting a trialling organisation with the development of their safety case will help ensure the safety case is appropriate for the environment they wish to operate in and will ultimately assist with reviewing the final safety case documentation.

EXAMPLES AND TEMPLATES

Safety case acceptance process	Outline of pre-defined process or requirements developed by the reviewer
Test facilities	Test facilities availability Test routes and tracks Vehicle monitoring capabilities
Local knowledge	Local road works Typical local operating conditions School locations and entry and exit times Bus routes Congestion at different times of the day Known accident blackspots
Pre-requisites	Policy statements
Insurance	Insurance requirements
Stakeholder engagement	Connections to the wider stakeholders within the test area such as local businesses, schools, residents, authorities to assist in stakeholder engagement

Table 2.1: Example information that a reviewer could provide to a trialling organisation prior to safety case development

2.3 Guidance on understanding operational safety case requirements

To assist with the review of a CAM technology trial safety case it is important that reviewers understand the recommended areas of focus for operational safety case development and understand the driving factors for safety case requirements.

When reviewing an operational safety case, there are three main risk factors that should be addressed by the safety case; the trial environment, the safety operator and the vehicle/system, as shown in Figure 2.4. The depth of information provided within each of these areas should be proportionate to the trial complexity.

The information included within an operational safety case is driven by the level of risk posed to all affected parties and the control over the three main risk factors. The operational risk assessment will require a balance between providing evidence of control and mitigations for additional risks where sufficient control cannot be achieved within these three areas.

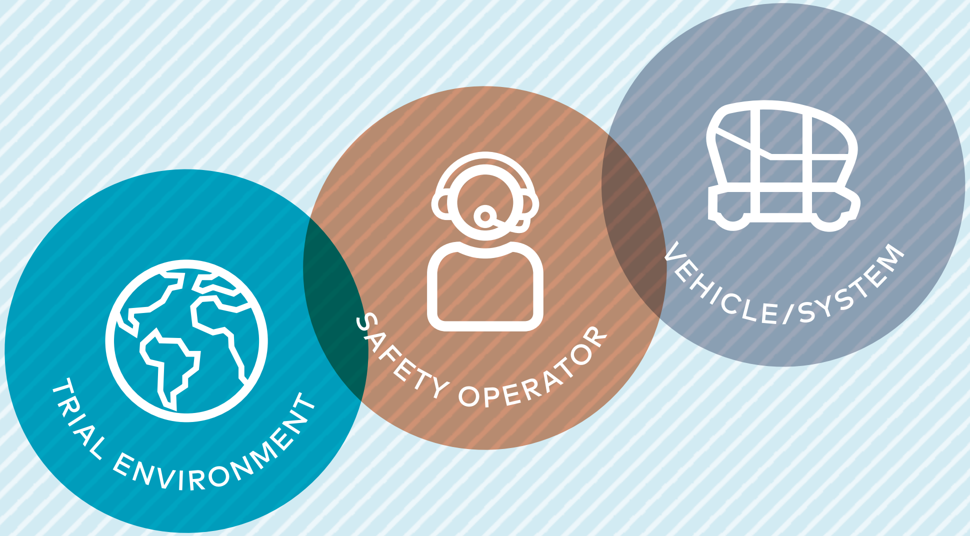

Figure 2.4: Safety Case Risk Factors

The level of control over each of the three risk factors will drive the balance between **evidence** and **risk mitigation** required within the operational safety case:

- where full control of risk factors can be achieved, the operational safety case should focus on providing evidence of such control
- where no or limited control of risk factors can be achieved, the safety case should provide risk mitigations within this area
- where partial control of a risk factor can be obtained, the safety case should provide a balance between evidencing control and mitigating the remaining risks.

As the level of control increases, risks can be demonstrated to be ALARP through evidence of that control rather than mitigations for lack of it, as shown in Figure 2.5.

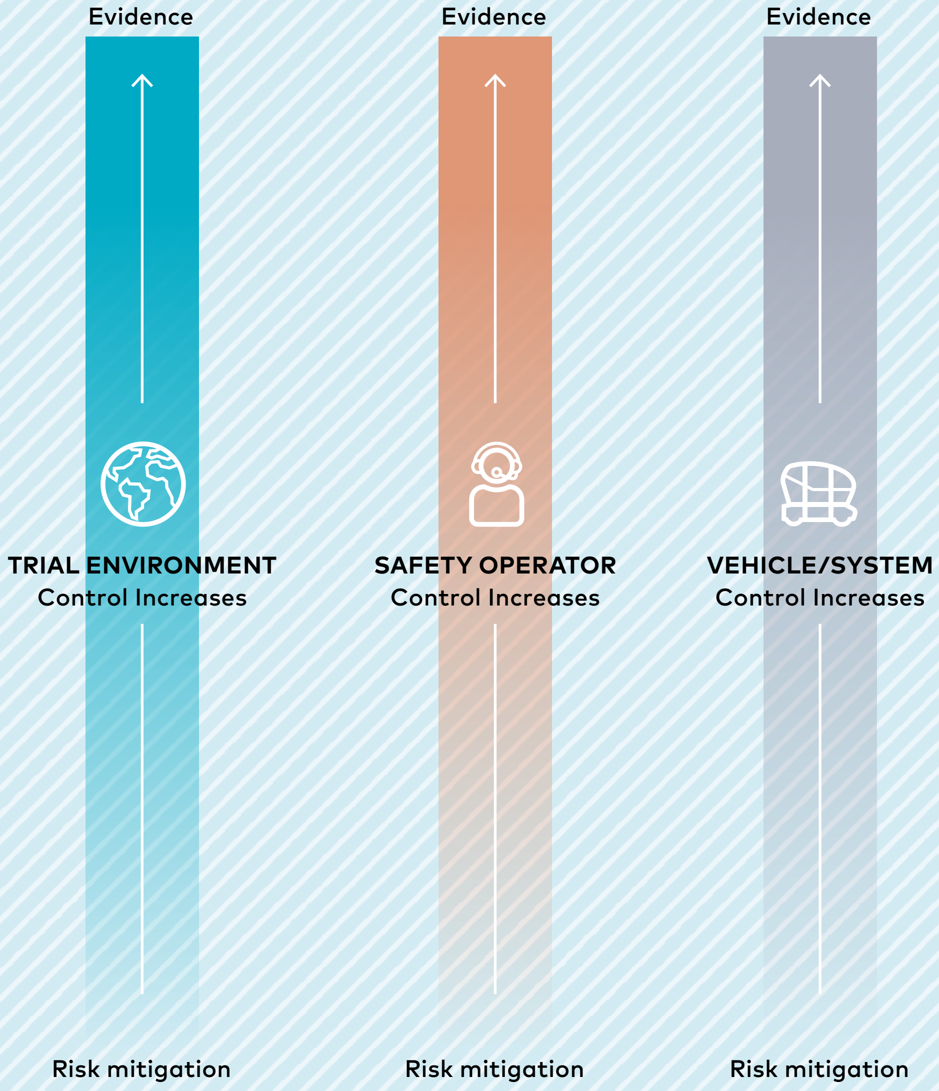

Figure 2.5: Safety case balance between evidence of safety and risk mitigation

KEY REVIEWER POINT

Reviewers should expect to see consideration of the level of control of the three risk factors (trial environment, safety operator and vehicle/system) and consideration of the resulting need to provide evidence of a high level of control and/or mitigation for a lack of control.

2.3.1 Determining the complexity of the trial

Trial complexity is a major factor in driving the level of detail required when providing evidence of control or details of identified risk mitigations. Low complexity trials will require a reduced level of detail within the operational safety case when compared to a high complexity trial.

Trial complexity is driven by trial design and operating conditions, and is directly related to the level of risk posed to affected parties during trial operation. The level of complexity therefore depends on how likely it is that an undesired event will occur and the consequence if such an event does occur.

There is a strong link between trial complexity and the Operational Design Domain (ODD, as detailed in Section 3). Factors to consider when determining trial complexity may include:

- Vehicle mass
- Traffic levels
- Proximity to vulnerable road users
- Weather
- Route features (e.g. proximity to schools)
- Temporary road structures (e.g. road works)
- Illumination levels
- Vehicle speed
- Road layout
- Junction types
- Road surface
- Presence of passengers
- Trial route length
- Special structure (e.g. tunnels, bridges)
- Vehicle operation and control (e.g. remote operation)

EXAMPLES AND TEMPLATES

Example factors for high complexity and low complexity trials are:

A high complexity trial may involve:

- Test scenarios that challenge the boundaries of the ODD
- High vehicle mass
- High vehicle speeds
- Busy high street environment with vulnerable road users
- Tall buildings affecting wireless connectivity
- Passengers carried in test vehicle
- Allowed to operate in heavy rain
- Remote operation

- **A low complexity trial may involve:**
- Trial design ensures that the boundaries of the defined ODD will not be challenged
- Safety driver with standard vehicle controls
- ADS installed within a standard production vehicle
- Low speed trial
- Test area with sparse traffic and no vulnerable road users
- Operation only in dry weather

As the complexity of a trial is driven by the trial design and operating conditions, methods for calculating complexity will vary between types of trials. A common approach has therefore not been defined for determining trial complexity, which should instead be assessed on a case-by-case basis using professional judgement. Trial complexity should not be equated to numeric values or be given a definite 'value'.

When testing in a proving ground (highly controlled environment) or operating a standard production test vehicle in manual mode (high level of operator control), the main safety case focus will be operational safety. Due to the low complexity and the need to adopt a proportionate approach, the safety case would typically consist solely of a risk assessment and a method statement. However, should such trials then progress to more complex environments or control levels, the scope and detail of the safety case would need to expand accordingly.

KEY REVIEWER POINT

Reviewers should expect to see evidence that the inherent complexity has been considered, based upon a review of the trial characteristics, and that the resulting safety case is of a level of detail that is proportionate to the complexity.

EXAMPLES AND TEMPLATES

The following examples may help reviewers get a feel for what to expect for different types of trials:

EXAMPLE **01**

An early development ADS integrated on a standard production vehicle being tested at low speed in a proving ground environment would be deemed a low complexity trial. Little evidence of system safety would be available due to the lack of prior testing, and the main risk mitigation would be that testing is being conducted in a very controlled environment. Details such as route assessments, ODD and security considerations may be minimal due to the controlled testing environment.

EXAMPLE **02**

A developed ADS operating in the public domain carrying passengers along a route through a busy urban environment would be deemed as a higher complexity trial. Significant detail would need to be provided within the safety case, demonstrating evidence of operator control and risk mitigations in place for the operating environment. Some evidence may be provided that the system performance is sufficient, although this may be deemed unnecessary due to the protection provided by operational safety measure.

EXAMPLE **03**

As per example 2, but with a remote safety operator who is able to supervise the vehicle and trigger an emergency stop. The safety operator is only able to take full control of the vehicle when performing low speed manoeuvring via a joystick, so this cannot be used for safety interventions. This would be a very high complexity trial, and in addition to the evidence and mitigations required in example 2, the trial would also need an extremely detailed system safety case and security case to demonstrate that the ADS could be trusted to operate without human intervention, and that the communications link for the remote emergency stop function is reliable and has adequate security.

2.4 Guidance on the support available to assist in the safety case review

The industry workshops conducted to inform this document indicated an ambition to set up a process to support Safety Case Creators, as it is recognised that a form of centralised support could aid all stakeholders in understanding what is required of them. This would provide a level of standardisation across CAM Testbed UK, thereby increasing the level of interoperability across testbeds.

As such a scheme is not currently available, trialling organisations should seek to work with the testbed(s) that they intend to trial at, from as early as possible within the safety case development process, to ensure agreement is reached upon a suitable approach.
It is hoped that the guidance contained within this document, and the corresponding one written for Safety Case Creators, will aid these discussions. In some cases, testbeds may be able to share information that will help with the Safety Case Creators, e.g. information on planned roadworks, known accident blackspots or details of the facilities available.

KEY REVIEWER POINT

Reviewers should ensure their expectations are consistent with the requirements in this document. Where possible, discussion between testbeds will help ensure consistency.

2.5 Summary of section 2

The level of detail contained within a trial safety case should be proportionate bearing in mind the inherent complexity of the proposed trial.

It should consider:

- How safe the level of control provided by the vehicle/system is
- Whether the surrounding trial environment can be controlled to manage safety
- The ability of the safety driver to ensure safety by making control inputs into the vehicle.

Where there is a high level of control of these factors, the safety case should provide evidence.

Where there is a lower level of control, the safety case should set out mitigation strategies to compensate for this lack of control.

The safety case will typically comprise of multiple documents; the ‘safety argument’ describes how separate pieces of evidence support the overall goal that the trial is acceptably safe.

The safety case will typically include consideration of operational safety, system safety and security. However, trials for immature technology will typically use operational safety measures (in particular, a safety driver) to compensate for not being able to provide evidence of system safety.

3.0 Definition of the trial characteristics

This section examines how to define the ODD and the intended vehicle behaviour. Providing a clear definition of the system and its environment is a vital precursor to the risk assessment and risk management processes detailed in later sections.

3.1 Operational design domain

Operational safety measures (Section 4) depend upon a clear and common understanding of the ODD, as the hazards that a trial will be exposed to will vary significantly depending upon the surroundings in which the vehicle is deployed. For example, a trial that will operate only upon motorways will have a different set of hazards to one taking place within an urban environment. Furthermore, it is essential that all personnel involved in maintaining operational safety during the trials (e.g. safety driver, test engineer) are familiar with the ODD so that they can recognise when the trial strays outside the ODD (e.g. due to unfavourable weather or the presence of a type of road user the system is not designed to react to).

Whereas safety operators are able to make inferences where a situation is not explicitly covered within the ODD definition and therefore not explicitly in or out of the ODD, based upon the inclusion or exclusion of similar permutations, an ADS has no such ability. Therefore, it would typically be necessary to specify the ODD in more detail where the safety case depends primarily upon assurance of system safety rather than upon operational safety measures.

A trial will have a finite ODD, whereas the real world has infinite variation. Therefore, there will always be a risk that the ADS may experience an ODD excursion. As such, the safety case must include consideration of an appropriate response in the event of an imminent ODD excursion, such as the ADS performing a Minimal Risk Manoeuvre (e.g. stopping in lane or navigating to a safe location to stop) or the Safety Driver taking manual control of the vehicle.

KEY REVIEWER POINT

Safety Case Reviewers therefore need to be satisfied that:

- the safety case documentation includes a description of the ODD
- the level of detail is proportionate to the needs of the trial
- the ODD specified for the trial is compatible with the testbed
- the safety case considers appropriate responses to prevent or mitigate excursions from the defined ODD.

ADVANCED TRIALS

3.2 Importance of the ODD where system safety is paramount

While defining an ODD is an integral aspect of the operational safety case, it is also the first step in the system safety process development. The inherent complexity within the real world means that an extremely large volume of testing is required to demonstrate that a system is able to operate safely without supervision. For many trials, this is circumvented by applying operational safety measures (e.g. the use of a safety driver or a controlled environment). However, for trials without such operational safety measures, an extensive test programme would need to be performed to demonstrate that the system safety is acceptable.

KEY REVIEWER POINT

Reviewers of such cases should satisfy themselves that there is sufficient evidence that:

a. the system has been tested in a sufficiently wide range of scenarios to provide acceptable coverage of the entire operational design domain, and

b. the system has shown an acceptable level of performance in those tests.

If these criteria are not met, trialling could still proceed if the trial ODD is revised to remove permutations where sufficient assurance has not been provided. Note that the amount of data produced by a large test programme would mean it would be disproportionate for reviewers to audit every test case undertaken, and therefore it is advised that a balanced approach should be taken such that reviewers can be reasonably satisfied that due diligence has been shown in providing coverage of the ODD, while ultimate responsibility for ensuring safety remains with the trialling organisation.

3.3 Relationship between ODD, desired behaviour and scenarios

When reviewing the ODD definition, it is important to check for any special ADS or other actor behaviour that might be relevant to the particular testbed; certain aspects of behaviour (both of the ADS and other actors around the ADS) may be restricted by the nature of the location. For example, the manoeuvre 'lane change' may be prohibited on a single lane undivided road. Once a trial route has been presented, reviewers should not only check for the ODD compliance with the trial route, but also the expected behaviour from the ADS and other actors on the trial route.

The ODD defines the range operating conditions an ADS could be exposed to and the behavioural competencies define the range of behaviours that the vehicle is capable of providing. Scenarios represent a combination of specific permutations of expected behaviour from an ADS with specific permutations from within the ODD. For example, performing an unprotected right turn (behaviour) on a single lane undivided crossroad (ODD) with oncoming traffic (behaviour) would be a scenario (illustrated in Figure 3.1). As there are typically many behaviour permutations that the vehicle is capable of and many ODD permutations that the vehicle could be exposed to, there will be many scenarios that are possible.

KEY REVIEWER POINT

Reviewers should be satisfied that the specified ODD is appropriate to the scenarios that will be tested, including consideration of both the ODD and the behavioural competencies expected of the ADS.

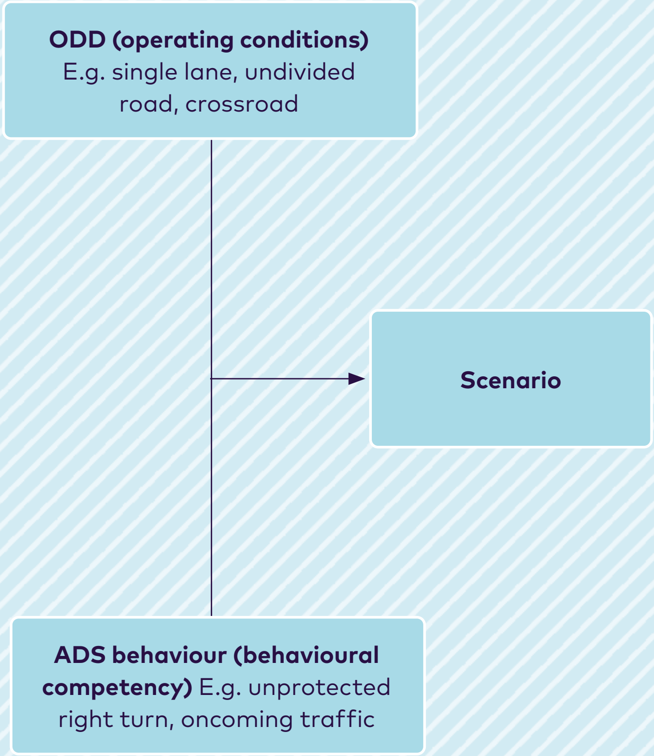
ODD (operating conditions) E.g. single lane, undivided road, crossroad
Scenario
ADS behaviour (behavioural competency) E.g. unprotected right turn, oncoming traffic

Figure 3.1: Relationship between ODD, behaviour and scenarios

3.4 Defining an ODD

While, depending on the perspective, the level of abstraction used within the ODD definition may vary, it is essential that the relevant stakeholders have an agreement on the ODD definition prior to the trial commencing. This agreement needs to take into consideration the level of detail appropriate to support other components of the safety case, e.g. operational risk assessment or system safety verification testing. It is recommended that Safety Case Creators use a standardised approach, such as that provided by the BSI PAS 1883 (2020), to define the ODD of the ADS, as use of a common taxonomy will ease the process of agreement between stakeholders.

As an ODD can be defined at levels of abstraction, reviewers should check whether level of detail within the ODD defined by the trialling organisation is compatible with the requirements of other stakeholders, e.g.:

- Testbed manager and/or operations team
- Local authorities
- Highway authorities
- Insurers.

KEY REVIEWER POINT

Reviewers should seek evidence of agreement between these stakeholders with respect to the ODD defined for the trial. In order to enable alignment, it is recommended that stakeholders use of a common taxonomy such as BSI PAS 1883.

BSI PAS 1883 provides a hierarchical taxonomy for ODD and features three top level attributes: 'scenery', 'environmental conditions' and 'dynamic elements'. Each of these are then further decomposed according to the desired level of abstraction deemed appropriate. For example, a trialling organisation may include junctions within their ODD. This would implicitly mean that the trial includes all types of junctions (roundabouts and intersections). However, if the ADS is only able to operate on certain types of junction, it would be necessary to decompose the junction attribute into different types of junction so that it can be defined which variants are with the ODD of the ADS and which are not. Reviewers should seek clarification from the Safety Case Creators on whether any attributes not explicitly referred to within the ODD definition are included within (permissive notation) or excluded from (restrictive notation) the defined trial ODD.

It is expected that Safety Case Creators may not use all the attributes within BSI PAS 1883 to define the trial ODD. Indeed, it is not a requirement that ODD definitions adhere to BSI PAS 1883, and any equivalent methodology that achieves the same overall objective should therefore be deemed acceptable.

The interdependence of the ODD attributes should be considered while defining the ODD of the trial. For example, for a particular trial, trialling organisations may define the top speed of the vehicle as 70 mph during the daytime on a sunny day, but may reduce the top speed to 40 mph on a rainy day. Such interdependence can be valuable to allow the system to be exposed to the broadest range of challenges possible in the given conditions.

KEY REVIEWER POINT

Reviewers should be satisfied that the definition of the ODD is sufficiently clear and unambiguous such that all stakeholders will have a common understanding. It should also be confirmed that the inclusions and exclusions are reasonable; for example, it would not be reasonable to exclude things that can occur unexpectedly, such as a horse and rider or an emergency vehicle upon a public road, unless some practical means to respond to the occurrence is incorporated into the safety case (such as a safety driver taking manual control).

In any instances where there is a discrepancy between the trial route and the trial ODD, reviewers should confirm the existence of an appropriate measure to ensure the trial remains within the ODD, such as the safety driver taking control for portions of the route that are incompatible.

3.5 ODD awareness

For the safe operation of the trial, it is essential that the personnel involved in the trial and/or the ADS itself are able to monitor the attributes referred to in the ODD in order to detect ODD excursions. Safety Case Reviewers should seek details of the monitoring/detecting mechanisms implemented; it is possible that the monitoring responsibilities for the ODD attributes to be split between:

- on board sensing (using sensors fitted to the ADS-equipped vehicle)
- off-board sensing (e.g. roadside weather station for information on weather attributes or a traffic management system for information on dynamic elements attribute), or
- safety operator judgement (by visual monitoring).

Where it has been deemed that an ODD excursion has occurred, or is about to occur, it would not be appropriate for the trial to continue. However, it should be borne in mind that the ODD definition can be written to allow a level of flexibility in when an attribute is in or out of scope by making use of interdependency between attributes (described in Section 3.4). For example, the ODD could specify no more than light rain for operation in the vicinity of other vehicles, but have an allowance for testing in heavy rain where there are no other vehicles nearby. This allows the capability of the system to be explored whilst also ensuring the trial remains within the specification that was considered during the risk assessment.

KEY REVIEWER POINT

In all cases, it should be confirmed that the ODD provides sufficient clarity for judgements to be made. In general, more detailed and quantitative definitions would be needed where the system is required to monitor the ODD and make decisions, in comparison to where this is done by humans.

3.6 Test scenarios

Specification of test scenarios is important for two reasons:

01

To provide relevant stakeholders with an understanding of what scenarios will be undertaken in the upcoming trial, which will aid understanding of the potential hazards within the operational safety analysis (Section 4), and

02

To provide stakeholders with data on what testing has previously been undertaken, which will be necessary where such evidence is needed due to the safety case being dependent upon assurance of system safety

The former reason is applicable to all trials, but will typically only require a high-level summary sufficient to support the risk assessment and method statement. The latter reason would only be applicable to advanced trials that do not have a safety driver present in the vehicle, and would require a large volume of detailed evidence.

A test scenario would typically comprise of:

- ODD elements present (e.g. oncoming car, pedestrian crossing road, rain)
- behaviour of the vehicle under test (e.g. required to enter roundabout and take second exit)
- pass criteria.

KEY REVIEWER POINT

Reviewers should ensure that the safety case makes clear what test cases will be undertaken in the trial, and where required as safety evidence, what test cases have already been undertaken.

The safety case would typically also include test cases that do not take the form of a scenario, to confirm the correct operation of various subsystems or components. For example, all trials that rely upon safety driver overrides would need test evidence to confirm that the override mechanisms (e.g. emergency cut-out button) work correctly, but aspects such as scenery or environmental conditions would not be a factor within these tests.

3.7 Summary of section 3

The safety case should include a definition of the ODD, such that all relevant stakeholders understand what surrounding features and characteristics are in scope for the trial. This information is needed to facilitate further safety analysis such as the risk assessment.

For the same reason, the behavioural competencies, i.e. the functionality that the vehicle is able to provide, should also be defined.

A scenario will combine a permutation that is possible within the bounds of the ODD with a permutation that is possible within the bounds of the behavioural competencies. All scenarios planned for the trial should therefore be consistent with the specified ODD and behavioural competencies.

Where evidence of system safety is needed to support the safety case (i.e. where a safety operator cannot be relied upon to correct errors by the system), the test cases that have been undertaken to verify the system safety should provide coverage of the full range of ODD and behavioural competencies that is possible.

4.0 The operational safety case

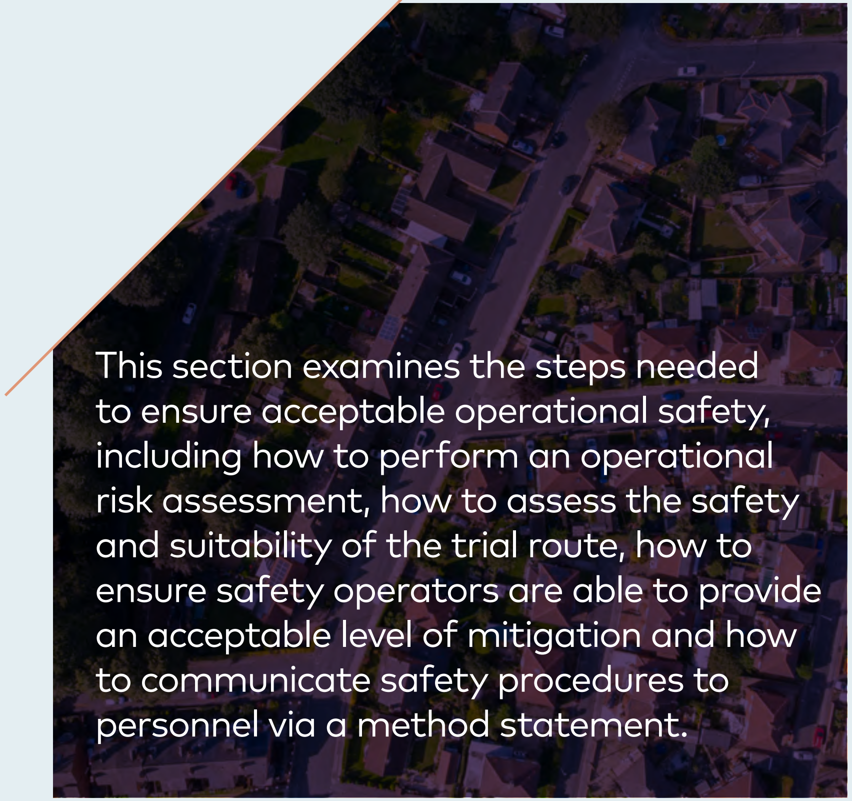

This section examines the steps needed to ensure acceptable operational safety, including how to perform an operational risk assessment, how to assess the safety and suitability of the trial route, how to ensure safety operators are able to provide an acceptable level of mitigation and how to communicate safety procedures to personnel via a method statement.

PRELIMINARY TRIALS

4.1 Operational risk assessment

4.1.1 What should be included in an operational risk assessment

An operational risk assessment forms the backbone of an operational safety case, and therefore Safety Case Reviewers should ensure that it has been developed thoroughly and is specific to the following three areas:

01
The trial **environment** (including the testbed and any other users within it)

02
The **safety driver or operator**

03
The **vehicle/system**

The development process used to create the operational risk assessment should also be described, including use of any industry standards and also any reliance upon pre-existing testbed risk assessments, hazard logs and associated templates. This description will provide the reviewer with an understanding of the resources used to develop the operational risk assessment, to satisfy them as to whether or not the Safety Case Creator has applied a proportionate level of rigor in their assessment of the risks posed by the trial.

The operational risk assessment should have been carried out by competent and informed individuals who are either directly or indirectly involved in the trials and have carried out operational risk assessments for similar trial activities before.

KEY REVIEWER POINT

Reviewers should ensure that the Safety Case Creator has an adequate process in place to allow sufficient consideration of all aspects of operational safety by suitably experienced and informed personnel.

4.1.2 Use of appropriate methodology and structure

The operational risk assessment method used to identify hazards and assess the resulting risks that could arise should be defined and justified by the Safety Case Creator. In the short to medium term, it is anticipated that trialling organisations will not have sufficient data to support quantitative risk assessments, therefore qualitative assessments should be expected and based on expert judgement (informed by available data). Any technical input received from experts to identify, assess and evaluate hazards should also be outlined in the operational risk assessment.

The high-level headings listed in the example below provide reviewers with a reference point, but it is acknowledged that some of the headings could vary in detail or sequence depending on the rigour required for the risk assessment; the reviewing process should remain flexible such that trialling organisations can use their own risk assessment formats, provided they achieve the underlying aim. Regardless of the headings used, they should be appropriate and understandable for the operational Safety Case Reviewer and for any other stakeholders required to use or update the operational risk assessment during the lifecycle of the trial.

EXAMPLES AND TEMPLATES

The structure of an operational risk assessment should include the following high-level headings:

- Hazards or hazardous scenarios identified
- Causes of those identified hazards
- Stakeholders affected by those hazards*
- Mitigations in place to control those hazards
- The likelihood of those hazards being realised
- The consequence severity of those hazards being realised
- The risk level based on an assessment of the likelihood combined with consequence severity

**Includes other testbed users, testbed staff, road users, members of the public and the emergency services.*

KEY REVIEWER POINT

Reviewers should be satisfied that the format used for the risk assessment allows hazards to be logged and a suitable system for rating the corresponding risks to be applied. Mitigations should also be logged where applicable. Flexibility should be allowed for trialling organisations to use their own in-house format, and therefore reviewers should not insist upon strict adherence to any particular set of column headings.

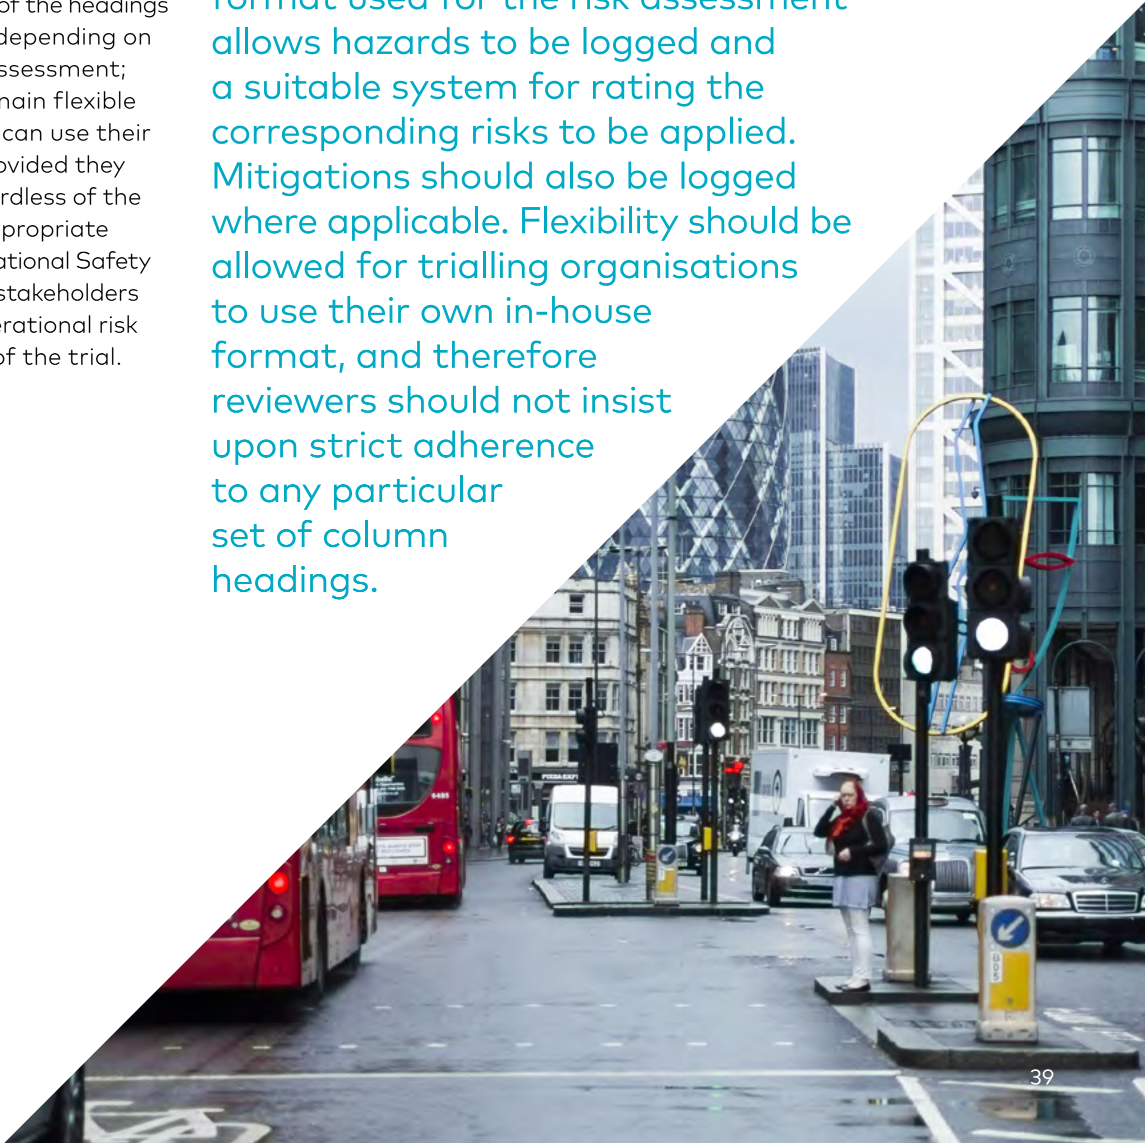

EXAMPLES AND TEMPLATES

	5 X 5 MATRIX	CONSEQUENCE SEVERITY				
		VERY LOW	LOW	MEDIUM	HIGH	VERY HIGH
LIKELIHOOD	Frequent	🟨	🟨	🟥	⬛	⬛
	Highly likely	🟩	🟨	🟥	⬛	⬛
	Likely	🟩	🟨	🟨	🟥	⬛
	Unlikely	🟩	🟩	🟨	🟥	⬛
	Improbable	🟩	🟩	🟨	🟥	🟥

Table 4.1: Example of a 5x5 qualitative risk matrix

4.1.3 Use of appropriate risk matrices

Determining the level of risk posed, understanding risk tolerability and prioritising risks and mitigations are key objectives within any risk assessment. Risk matrices are a useful tool for quickly defining risk levels by assessing the likelihood of a risk being realised with the consequence severity of that risk. Qualitative or quantitative matrices can be used and, like the risk assessment methodology, should be determined based on the availability of sufficient data to support any assumptions.

An example 5x5 qualitative risk matrix is shown in Table 4.1, which has been adapted from the guidelines outlined in BS ISO 31000 (2018). A 5x5 risk matrix is one of the most common risk matrix formats, although a 3x3 method could be appropriate where there are fewer, less complex risks posed by the activity being assessed. Testbeds may be able to assist Safety Case Creators by providing a risk matrix template to a pre-existing format or by referring to a matrix in a standard. However, use of these templates should be optional, and reviewers should not impose a requirement to use a specific format, provided that the one adopted follows risk management good practice and has suitable tolerability levels, and reasonable justification for its selection has been given. Safety Case Reviewers should be satisfied that the matrix used is proportionate to the complexity of the trial and to the trial environment.

KEY REVIEWER POINT

Reviewers should confirm that an appropriate methodology for risk assessment has been developed that enables risks to be rated and prioritised such that the need for further mitigation can be identified. Reviewers should accept any approach that achieves this underlying objective and applies suitably professional practice; trialling organisations should be given the freedom to apply their own in-house risk matrices.

4.1.4 Consideration of relevant hazardous scenarios and mitigations

Sources of hazardous scenarios, such as the base vehicle, ADS, safety operator, route, infrastructure, and external dependencies, are outlined in BSI PAS 1881. Safety Case Reviewers should familiarise themselves with these sources and ensure that the sources referenced by the Safety Case Creator are appropriate, and are as exhaustive as possible based on the risks posed by the trial and trial environment.

Safety Case Reviewers should also take into consideration the novelty and maturity of the vehicle and ACS being tested, as this could influence the types of hazardous scenarios identified, depending upon the trial environment. Attributes such as unique vehicle appearance (in terms of shape, size or features), distinctive signals, or different ways in which the vehicle behaves and can be manoeuvred could result in additional hazards.

EXAMPLES AND TEMPLATES

Key hazards to consider when trialling in public road environments where vulnerable road users of various types could be present include, but are not limited to, the following examples:

Collision with pedestrian caused by:

- pedestrian stepping into the road (distraction, failing to look, inaudible or inconspicuous trial or test vehicle)
- pedestrian testing the response of the trial vehicle
- trial vehicle diverges from intended path
- unexpected acceleration of trial or test vehicle.

Collision with cyclist caused by:

- cyclist close following the trial or test vehicle
- cyclist testing the response of the trial vehicle
- cyclist unexpectedly swerving or changing direction
- trial vehicle unable to stop
- Obscured or inconspicuous cyclist.

Collision with motorcyclist caused by:

- motorcyclist close following the trial or test vehicle
- motorcyclist weaving through traffic
- obscured or inconspicuous motorcyclist.

Collision with horse and or horse rider caused by:

- trial vehicle speeding on approach to the horse rider or when overtaking
- unexpected horse or animal behaviour
- obscured or inconspicuous horse or horse rider.

Identification of hazardous scenarios should include interactions with third parties and all types of foreseeable road users in the given trial environment. In circumstances where trials are undertaken on public roads, Safety Case Creators should provide particular focus on interactions with vulnerable road users, defined by Rule 204 of the Highway Code (2019) as pedestrians, cyclists, motorcyclists and horse riders. Rule 204 also states awareness should be given to children, older and disabled people and learner and novice drivers and riders who could be considered to be vulnerable; these should also be considered and evaluated by the Safety Case Reviewer.

Further to Rules 219 to 225 of the Highway Code (2019), interactions with other vehicles, such as heavy goods vehicles (HGVs), light goods vehicles (LGVs), buses, coaches, recovery vehicles and emergency and incident support vehicles, could also introduce hazardous scenarios.

The causes and consequences of the identified hazardous scenarios should be used to identify suitable mitigations.

Examples of such mitigations include producing operational guidance for the appropriate personnel (e.g. safety driver or test engineer), introducing operational controls (such as the safety driver resuming manual control pre-emptively) or refining the operational design domain (ODD). Table 4.2 provides an example of how two hazards within a risk assessment could be logged, risk assessed and mitigated.

EXAMPLES AND TEMPLATES

Some of the hazards resulting from other vehicles types are:

Collision with HGV or LGV caused by:

- trial vehicle or safety driver or operator misjudges the turning circle and size of the HGV or LGV, relative to the road space available
- trial vehicle closely following an HGV or LGV, obscuring the road, road features and hazards ahead.

Collision with bus or coach caused by:

- trial vehicle closely following a bus or coach
- trial vehicle or safety operator failing to allow the bus or coach to pull away from stops.

Collision with emergency or incident support vehicle caused by:

- stopping or pulling the trial or test vehicle over in a hazardous location
- safety driver or operator failing to observe or anticipate an approaching emergency or incident support vehicle.

EXAMPLES AND TEMPLATES

HAZARDOUS SCENARIO	CAUSES	PARTIES AFFECTED	MITIGATIONS	LIKELIHOOD	SEVERITY	RISK LEVEL
Collision with pedestrian	Pedestrian stepping into the road Pedestrian testing the trial or test vehicle Trial or test vehicle diverges from intended path Unexpected acceleration of trial or test vehicle	Pedestrians Safety driver Test assistant Members of the public	Conspicuous and audible test or trial vehicle Fully trained safety driver Fully tested ADS and object detection capability Public awareness campaign to make members of the public aware of the tests or trials	Unlikely	High	High
Collision with cyclist	Cyclist close following the trial or test vehicle Cyclist testing the trial or test vehicle Cyclist unexpectedly swerving or changing direction Trial or test vehicle unable to stop Obscured or inconspicuous cyclist	Cyclists Safety driver Test assistant Members of the public	Conspicuous and audible test or trial vehicle Fully trained safety driver Fully tested ADS and object detection capability Public awareness campaign to make members of the public aware of the tests or trials	Unlikely	High	High

Table 4.2: Example of how hazards within a risk assessment could be logged, prioritised and mitigated

KEY REVIEWER POINT

Safety Case Reviewers should ensure that the risk presented by each hazard has been assessed, and the causes and consequences of each hazard used to establish mitigations. Mitigations should be commensurate with the risks posed and could comprise system or operational controls.

4.1.5 Consideration of input from stakeholders

Safety Case Creators should describe the input received from stakeholders during the development of the operational risk assessment, both during hazard identification and review of the operational risk assessment where applicable. The rationale for doing so is to provide the reviewer with assurance that all local knowledge and foreseeable hazards associated with the trial are documented and evaluated.

Further information on who might be a stakeholder and what form on contact should take place can be found in Section 7.

4.1.6 What should be included in monitoring, reporting and continuous improvement?

An operational risk assessment and operational safety case are live documents and should be reviewed and updated periodically, or immediately if required following an undesired event or significant system, environment or process change. Such required updates should take place prior to further trialling. A monitoring and analysis plan should be in place which demonstrates that the key hazards and assumptions made in the risk assessment are being monitored throughout trials. Similarly, the data and feedback being captured, and the mechanisms being used to monitor and capture it, should be clearly defined and communicated to the relevant parties.

Guidance on what Safety Case Reviewers should look for to confirm the existence of a suitable process for updating the operational safety case can be found in Section 7.

PRELIMINARY TRIALS

4.2 Operational guidance

4.2.1 Relationship between operational guidance and operational risk assessment

The operational risk assessment will inform what should be included in operational guidance regarding control measures and mitigations. The operational guidance should be proportionate to the risks associated with the trial, so Safety Case Reviewers should ensure that operational mitigations mirror the proposed measures identified through the risk assessment process. Where this is not the case, justification should be provided, e.g. through application of alternative control measures.

The causes and consequences of the hazardous scenarios identified in the operational risk assessment should be used to identify suitable mitigations. Again, this will differ by trial activity and environment but could include producing tailored operational guidance that is proportionate to the environment, safety operator, vehicle and activity.

4.2.2 What operational guidance should reviewers look for?

As described in Section 2.1, operational guidance will differ between trials depending on the environment, safety driver or operator and vehicle/system. In terms of the types of guidance to consider, BSI PAS 1881 (2020) describes a comprehensive list which has been summarised below:

- Method statement
- Abort policy or procedure
- Safe operation of the ADS on the given route(s)
- Safety driver or operator policies
- Vehicle storage and security
- Vehicle maintenance, inspection and cleaning procedures
- Vehicle fuelling and charging
- Vehicle recovery plan
- Incident reporting policy or procedure
- Emergency response plan and crisis communication plan

However, the following essential operational guidance should be provided for all safety cases and in all testbed environments:

- Method statement
- Emergency response plan and crisis communication plan
- Incident reporting procedure

Note that in some cases testbeds will provide the latter two documents. If this is the case, it may not be necessary for trialling organisations to produce their own versions, although they may choose to in order to instigate additional policies and procedures that are supplementary to those of the testbed. Each of these documents are covered in more detail within the following two sections.

4.2.2.1 Method statement

Safety Case Creators should develop a method statement, which describes the sequence of tasks being undertaken for the trial activities and how they are being carried out in a safe manner. Safety Case Reviewers should ensure that the method statement describes, in a level of detail consistent with the complexity of trial and environment, the following:

- The roles and responsibilities of the trial team
- An overview of the key risks and mitigations
- Reference to appropriate operational guidance documents
- An outline of the planned trial schedule and objectives

Testbeds may have existing method statement templates that can be provided to operational Safety Case Creators and reviewers, although safety cases should not be restricted to any one specific format or structure. The method statement should be shared with all personnel and stakeholders that have a safety-related role within the trials, or within the wider project. An example of a possible structure for the method statement is provided in Table 4.3.

EXAMPLES AND TEMPLATES

EXAMPLE SECTION HEADINGS	EXAMPLE SECTION CONTENT
Aims and objectives of the test or trial	Overview of test or trial aims and objectives
Sequence of tasks being undertaken	Description of what is being undertaken, when, and by whom Description of how the tasks are being carried out in a safe manner and the associated roles and responsibilities of the test or trial team
Overview of key risks and mitigations	Outline of the key risks and mitigations for the test or trial Reference to relevant operational guidance documents, including incident reporting
Route and test or trial schedule	Overview of the route being used Outline of the operating dates and times, including breaks
Key points of contact	Key points of contact for the lead test or trial organisation and any other relevant parties involved in the activities

Table 4.3: Example of what content could be included within a method statement

KEY REVIEWER POINT

Reviewers should expect to see a method statement for all trials, regardless of the technology or the ODD. This should be sufficient for all personnel involved in the trial to understand what is expected of them in order to ensure safety.

4.2.2.2 Contingency planning: emergency response plan and crisis communications plan

An emergency response plan and accompanying crisis communications plan should be in place for all trial activities to ensure responses to incidents and emergencies are dealt with in an efficient and effective way. These documents should define the actions required and roles and responsibilities of the trials team in the event of a serious incident, to protect against further harm to individuals, financial loss or reputational damage. Where applicable, consideration should be given to the needs of affected stakeholders, e.g. emergency services, landowners.

Where testbeds have pre-existing emergency response and crisis communications plans, reviewers should ensure that the operational safety case is consistent with these. Bespoke contingency plans may be required for trials in other environments; this is described in further detail in BSI PAS 1881 (2020).

KEY REVIEWER POINT

Reviewers should expect to see an emergency response plan and a crisis communications plan if one is not provided by the testbed. Optionally, these two plans could be covered within a single document.

4.2.2.3 Incident reporting procedure

The means for reporting and analysing incidents should be described in an incident reporting procedure and the operational safety case. This should also be replicated, or referenced, in the emergency response plan. Further guidance on incident reporting procedures can be found in Section 7.2.1.2.

KEY REVIEWER POINT

Reviewers should be satisfied that trialling organisations are familiar with any incident reporting procedure that the testbed may have. If no such procedure exists, reviewers should ensure that one is provided within the safety case.

PRELIMINARY TRIALS

4.3 Route selection and assessment

4.3.1 Evidencing route selection criteria and assessment methodology

The route selected for trials should be appropriate for the capabilities of the vehicle and the ADS, the trial scenarios, and also the ODD where this has already been defined. Where the trial scenario pushes the boundaries of the defined ODD, additional consideration should be given to the level of control over the trial environment, the space available around the vehicle (to ensure there is sufficient room to allow the safety operator to react before a collision occurs), the capability of the ADS and whether new hazards could arise or existing ones become intolerable.

Route selection criteria will vary depending on these factors and upon the complexity and controllability of the trial environment. Testbeds may be able to assist by providing trialling organisations with information on what to expect on the routes available, and therefore the ODDs they support, to aid their route selection and ODD definition where applicable. The following information should be considered as minimum criteria when selecting a route in all types of trial environment:

- Controllability of the environment (for example, is exclusive use of the facility available?)
- Space available (length and width) relative to the vehicle type, size and capabilities and ACS capabilities
- Carriageway type
- Posted speed limit
- Presence of road signs or markings
- Hazardous (or potentially hazardous) track or road features
- Presence and types of track or road features (e.g. junctions, roundabouts, slip roads) and any collision 'hot-spots'
- Track or road geometry and topography
- Track or road surface condition
- Direction of travel or traffic

In addition to the criteria above, when selecting a trial route in a public or private road environment, trialling organisations should consider road user compositions and characteristics, and also the traffic flow. Where any of the listed information above is unavailable at the point of route selection, it should be obtained during the subsequent route assessment. Other factors, such as the land use adjacent to the road, should also be assessed where it has the potential to influence road user compositions and flows.

Safety Case Reviewers should confirm that the route assessment method used and the conclusions from that assessment are appropriately evidenced and justified. The route assessment methodology carried out should be commensurate with the following:

- The objectives of the trial scenario and ODD
- The level of controllability and control required within the environment
- The maturity of the ADS being used
- The presence of higher risk locations that could require additional safety assurance

The route assessment may have used a staged approach, using desktop and on-site assessments, or an on-site assessment on its own. Either approach could be acceptable and should be justified by the Safety Case Creator in the route safety assessment and operational safety case.

The reviewer should evaluate the assessments consideration of the following:

- Alignment of the route with the defined ODD
- Presence of static and dynamic hazards
- Presence of any triggering events.

Static hazards could be considered to be hazards that are permanently stationery, such as street furniture including bus stops, bridge abutments or zebra crossings. Dynamic hazards could be considered to be hazards that have the ability to move, such as an animal or a refuse collection vehicle - a parked car should be considered dynamic as there is a possibility for it to move at any given time. Triggering events are events that can trigger potentially hazardous behaviour; for example, inconspicuous road markings or road signs that could lead to illegal or poor road user behaviour or a collision.

Where ODD boundaries are being challenged, the route assessment should consider the requirement for additional route mitigations to control the risks posed, depending on the complexity and controllability of the trial environment. This could include requirements to monitor and control other traffic or limitations to the hours during which trials can take place.

The route assessment methodology and findings should be documented in the operational safety case to provide assurance that the selected route is compatible with the vehicle, the trial scenarios, and the ODD. Similarly, the route-specific mitigations and control measures used should be documented.

KEY REVIEWER POINT

Reviewers of safety cases should be satisfied that a proportionate process has been used to assess the safety and suitability of the route.

4.3.2 Relationship with the ODD definition

Safety Case Creators should have used route selection and assessment to either help define an ODD for a trial or to help validate a pre-defined ODD. As discussed in Section 4.3.1, the route selection criteria and assessment methodology for trials should also be appropriate for the ODD. Therefore, regardless of the trial environment, potential route hazards should be identified and understood relative to the ODD and trial objectives. Safety Case Reviewers may find the list of ODD attributes included within BSI PAS 1883 (2020) provides a valuable benchmark when validating the route assessment.

KEY REVIEWER POINT

Reviewers should confirm that a proportionate range of hazards have been identified, which matches the ODD. Depending upon the trial complexity and risk, this could take the form of a full audit or of sampling a proportion of hazards to get a feel for the rigour applied.

4.3.3 Relationship with the operational risk assessment

Route assessments will inform what needs to be considered in operational guidance regarding control measures and mitigations. Consequently, operational controls should correspond with the proposed measures identified through the risk assessment process. In particular, the causes and consequences of the hazardous scenarios identified in the operational risk assessment should be used to identify route-specific mitigations.

This could include producing tailored operational guidance that places control measures on how to interact with specific road users and road features, as well as identifying places of relative safety in the event of a breakdown or pause in tests. Furthermore, controls may be required to ensure that trials on the chosen route remain within the defined ODD. This could include monitoring weather or environmental conditions and initiating a suspension of trial activities where required.

KEY REVIEWER POINT

Reviewers should be satisfied that hazards identified within the route assessment have been captured within the risk assessment, and risks mitigated where appropriate.

4.3.4 Guidance on performance monitoring

Performance monitoring could involve monitoring the performance of the ADS, safety operator or dynamic hazards along the trial route. In particular, methods to monitor the safety operator performance, such as a trial engineer sat alongside them ensuring they are acting appropriately and do not appear to be drowsy, may help with justifying the ability to rely upon the safety driver to intervene safely at all times.

Evidence provided to Safety Case Reviewers should be described in the operational safety case and operational risk assessment with reference to any applicable operational guidance or monitoring plans. Requirements will likely be more detailed for performance monitoring during remotely operated trials, particularly concerning safety operator performance.

KEY REVIEWER POINT

Reviewers should be satisfied that the performance monitoring being undertaken during the trials is proportionate to the risk posed by the trial activity.

4.4 Safe operation and control

4.4.1 Relationship with the defined ODD

The safety case should include evidence that the safety operator has a very high level of familiarity and understanding of the ODD, to ensure that they understand when the vehicle may exceed the ODD and can intervene promptly and safely to manage the situation. It is important to note that an ODD is specific to a system, environment and trial scenario, and therefore a safety driver or operator must be familiar with all ODDs they may be operating under, and which one applies during any given test session.

KEY REVIEWER POINT

Safety Case Reviewers should be satisfied that a process exists to ensure that the safety operator has a thorough understanding of the ODD.

4.4.2 Evidencing an appropriate level of control and a minimal risk condition

The safety case should include documentation of how to achieve the minimal risk condition (MRC) within the trial. The MRC is defined as a stable condition to which a human driver or ADS brings a vehicle in order to minimise the risk of an undesired event. This may involve the ADS stopping the vehicle in the safest location available (avoiding stopping in a live lane if possible). There is some discrepancy within literature as to whether a non-stopped condition that minimises risk (such as a safety driver assuming control but keeping the vehicle in motion) is classed as an MRC; for example, BSI PAS 1881 (2020) does not restrict the MRM definition to 'stopped', whereas the BSI CAV Vocabulary (2020) does. Regardless of terminology, however, alternative options to achieving minimal risk, such as the system continuing in a 'safe mode' with reduced functionality or handover to manual control, should be considered and documented in the safety case.

As such, the definition of what constitutes an MRC should be particular to each trial and safety case. Optionally, there may be multiple MRCs defined for selection according to the prevailing situation. If a safety operator is in any way responsible for initiating or monitoring transition to an MRC, a process should be in place to ensure they are aware of all defined MRCs, how they are to be achieved, and under what circumstances. Where achieving an MRC requires a manoeuvre to be performed, that manoeuvre is referred to as the minimal risk manoeuvre (MRM).

The level of control, and therefore the level of evidence required, may vary considerably depending on the vehicle and system under test. For example, the interfaces used to control a pod may be fewer and more basic than for a conventional vehicle. The ability of the ADS to perform an MRM autonomously without driver input, in a safe and reliable manner, will depend upon the complexity and maturity of the technologies used (such as the sensors and the method of pathfinding).

Evidence should be provided to demonstrate a safe level of control, accomplishment of the MRC(s), and the required MRM(s) to reach the MRC(s). The selection of a suitable MRC will depend on the ODD and upon the triggering circumstances that make it necessary; for example on a high-speed road, pulling the vehicle over in a place of relative safety would be safer and more favourable than an e-stop that leaves the vehicle stationary in a live lane.

Required evidence will vary for different vehicles, use cases and trials but the following evidence should be included as a minimum:

- **Minimal risk condition** – the different mechanisms to achieve this, with evidence that it can be achieved in sufficient time to ensure that safety is maintained
- **Safety driver or operator human factors** – monitoring requirements and mitigations for safety drivers or operators; for example, ensuring the driver or operator has a suitable view of the vehicle to maintain situational awareness
- **Communications links** – the resilience and latency of the methods used and ways of maintaining these links, or using safe and suitable alternatives where relevant. This latter consideration is only applicable where there is a remote operator and hence the communications link is safety critical.

KEY REVIEWER POINT

Safety Case Reviewers should confirm that one or more MRCs or other means of minimising risk (e.g. handing over to an onboard safety driver) have been defined. These should cover all foreseeable permutations within the ODD. There should also be evidence that the level of control available to the safety operator is appropriate bearing in mind the nature of the trial environment.

PRELIMINARY TRIALS

4.4.3 Role of safety Drivers and Operators

Safety Case Creators should describe the role and responsibilities of the safety driver or operator in the operational safety case and any accompanying operational guidance should demonstrate that safety operators to have knowledge and competency sufficient to:

- understand and safely operate the ADS
- be aware of the specific vehicle capabilities and limitations
- understand the defined trial ODD, scenarios and objectives
- understand likely failure modes and how to mitigate them
- understand and comply with the applicable operational guidance
- understand the conditions that require the trial to be aborted and the procedures to follow in the event of a hazard or undesired event.

A member of the trial team should also be responsible for regularly checking all the control systems and driver or operator interfaces, and an indicative process for doing this should be outlined. This role could be assigned to the safety operator, but where there is concern that this could increase risk by distracting their attention, assignment of this task to another individual who is able to communicate directly with the safety driver should be considered as an alternative.

In addition, a plan for managing and mitigating against operator fatigue and distraction should be included. This could include maximum operating hours, regular breaks, a driver fatigue monitoring system, requiring another member of the test team to observe operator attention levels, and giving the operator the discretion to stop at any time if they do not feel safe to continue.

The safety case should provide an outline of what a safety operator needs to be capable of for the given trial and provide evidence to demonstrate that this can be achieved. This may include simulator or test track demonstrations where written evidence on its own is not enough, and will require steps to eliminate or mitigate the risk where it cannot be shown that the safety operator can intervene successfully.

EXAMPLES AND TEMPLATES

If a section of public road is so narrow that data assessing driver interventions via fault-injection testing on a proving ground suggests that the safety driver would not always be able to correct a steering error before the vehicle enters the path of oncoming traffic, the safety driver cannot be relied upon as a protection mechanism in that particular vehicle upon that particular stretch of road. Further mitigation would therefore be needed, such as removing the narrow section of road from the trial route, a safety driver taking manual control for that portion of the trial, or controlling the route such that oncoming vehicles will not be encountered.

KEY REVIEWER POINT

Safety Case Reviewers should require evidence to show that the safety operator is able to intervene reliably and safely whenever necessary and that consideration has been given to ensuring that they are able to maintain their performance throughout the trial.

PRELIMINARY TRIALS

4.4.4 Guidance on safety driver or operator training

Evidence should be provided regarding the level of training given and how it will mitigate against any undesired events in relation to the proposed activities. Safety driver or operator training should include maintaining control of the vehicle safely within and beyond the ODD, and the standard operating criteria for the vehicle.

Safety driver training should cover two main areas:

1. 'General' driving including familiarity with the relevant road traffic legislation and good driving practice, as described by the Highway Code (2019). This should also include development of skills in basic vehicle control and advanced driving skills such as maintaining control of the vehicle at the edge of its performance envelope.
2. 'Specific' driving related to the automated functionality of the trial vehicle, how it is designed to behave and operate, alerts and warnings, how to take control or manage a handover of control situation and any specific detail concerning the trial or route being used.

In many cases, trialling organisations will have their own in-house training to cover 'general' driving, and suitable courses are available commercially. However, 'specific' training would need to be provided as part of the safety processes for each particular trial. Specific training may include knowledge of the operational guidance, hazards and mitigations, defined ODD, handover points and any safety protocols (e.g. vehicle checks at the start and end of a trial day).

Ultimately, the evidence required to ensure adequate safety operator training has been carried out will vary depending on the level of control, the vehicle and trial environment. Safety Case Reviewers should use this guidance to make informed decisions as to whether the standard of training followed is appropriate.

KEY REVIEWER POINT

Safety Case Reviewers should confirm that a suitable process is in place to ensure that drivers are suitably trained and assessed such that they have a high level of general competence and a detailed knowledge of safety-related aspects of the particular trial and vehicle.

ADVANCED TRIALS

4.4.5 Remote safety operators

Remote safety operators should be trained to at least the same level required for in-vehicle safety drivers with additional focus on the following areas:

- How safe control of the vehicle is always maintained in the trial environment through the ability of the safety operator to make control inputs (e.g. via a games controller, joystick, or remote steering wheel and pedals)
- How the safety operator is alerted that action is required
- How network and communication links are made robust, how they are monitored and what fail-safes in place
- How the behaviour of the vehicle will be monitored remotely, with full visibility of the vehicle and surroundings provided such that situational awareness can be maintained

Furthermore, many of the areas covered elsewhere in this guidance would need far more detailed consideration in order to provide evidence that safe control can still be provided by the safety operator or to provide mitigation for the lack of ability to provide that safe control. For example, far more robust demonstration of the ability to perform a suitable MRM would be required relative to trials with a manual driver within a vehicle.

The ability to maintain safe control of the vehicle should be demonstrated through extensive safety testing, and potentially through demonstrations to testbeds or other stakeholders. An overview of the human machine interface (HMI) or user interface display and warnings should be described and be understandable to both the operator and other stakeholders.

KEY REVIEWER POINT

Safety Case Reviewers should require extensive test evidence demonstrating the robustness of all subsystems that support remote monitoring, and extensive analysis of the human factors involved (both in terms of feedback to the safety operator and the ability of the safety operator to provide safe control inputs).

4.5 Summary of section 4

The safety case should include an operational risk assessment that allows hazards and their mitigations to be logged and prioritised. This will be a key component of all safety cases.

The safety case should also include operational guidance such as a method statement and emergency response plan; such documents help convey safe working practices identified by the risk assessment to the relevant personnel in a clear and concise manner.

The trial route should be assessed to identify hazards and ensure compatibility with the ODD of the vehicle.

The safety case should include processes for ensuring safety operators are able to intervene safely. This includes training in general driving skills such as understanding the relevant rules and having suitable car control skills, and also specific training to ensure familiarity with the characteristics of the vehicle and trial.

5.0 The system safety case

This section examines what evidence of system safety may be required within a trial safety case. Guidance is provided on suitable methods to analyse the safety of systems, how to select appropriate test scenarios and how simulation can be used to provide safety evidence.

5.1 The need for system safety assurance

The assurance of system safety requires a body of test evidence to be collected in order to demonstrate that the likelihood of the system making a safety-critical error is sufficiently low for this residual risk to be acceptable. Because of the complexity of both the systems themselves and the environments in which they operate, this body of test evidence would need to be extremely large in order to provide coverage of the range of permutations that the vehicle could be exposed to. It is for this reason that many automated vehicle trials rely upon operational safety measures (such as a safety driver) to provide a protection mechanism such that system safety of the ADS performance does not need to be proven.

Where the overall safety case relies upon operational safety measures and therefore does not require the performance of the ADS to be verified, the only system safety needing to be evidenced would be the safe and reliable operation of any overrides required as part of the operational safety case.
In particular, any driver overrides (e.g. an emergency cut-out button to disable the ADS so the vehicle returns to manual driving mode or an emergency stop button to cause a low-speed vehicle to perform an emergency stop) would need to be robustly assured.

Such verification would not require 'scenario-based testing', as introduced in Section 5.2, and could be achieved using traditional systems engineering techniques such as:

- analysis methods such as an FMEA (Failure Modes and Effects Analysis) to confirm that the design is suitable for providing the functionality safely and robustly
- physical testing of the override to ensure that it is safe, reliable, fast and ergonomic

Where remote operation is used as a safety backup, the level of detail needed within this system safety analysis would be greater, as the robustness of a remote communication system would be more complex to assure. Nonetheless, this could still be achieved with standard good practice for systems engineering, safety engineering and cyber security.

KEY REVIEWER POINT

Safety Case Reviewers should seek evidence that any systems or subsystems required as part of an operational safety mitigation (in particular, safety operator overrides) have been subjected to rigorous design analysis and testing.

However, where it is necessary to ensure that the ADS performance itself is safe due to the safety operator having limited or no ability to intervene sufficiently, an extensive process will be required to generate an extensive set of test scenarios. These scenarios would need to provide acceptable assurance that the vehicle will perform safely in any situation that is reasonably foreseeable within the ODD, including rare permutations ('edge cases').

This would therefore require a test programme of a similar magnitude as that which would be required prior to the deployment of commercially available production automated vehicles. Sections 5.2 to 5.4 provide a summary of the key considerations for assuring system safety for such trials; for further information, readers may find UL 4600 (2020), SaFAD (2019) and RAND (2020) informative.

KEY REVIEWER POINT

If the safety argument depends upon the ADS being demonstrated to be able to function safely without recourse to a safety operator, reviewers should expect to see that the vehicle has been subjected to an extensive programme of testing to assure system safety. As such, safety cases that primarily rely upon system safety will typically need a more detailed review of the safety case and require a reviewer with more specialism domain knowledge, when compared to a safety case that primarily relies upon operational safety.

ADVANCED TRIALS

5.2 Scenario-based testing

Scenarios generation can be classified into two categories:

- **Data based scenario generation:** This involves analysing road collision databases and insurance claim records to identify parameters that contribute disproportionately to accidents. Additionally, such an analysis can provide insights into road collision hotspots which can then influence route selection for a trial.
- **Knowledge based scenario generation:** This involves analysis of the system architecture to identify potential failure modes and hazards, using established safety analysis methodologies such as STPA (Systems Theoretic Process Analysis), FMEA, FTA (Fault Tree Analysis) or HAZOP (Hazard and Operability Study).

KEY REVIEWER POINT

While reviewing the evidence for a safety case that relies upon scenario-based testing to demonstrate that the system is safe to operate without human intervention, reviewers should look for the use of one or both of the methods for test scenario generation. However, depending on the trial complexity, the number and detail of test scenarios may vary. Furthermore, reviewers should check that the test scenarios identified represent the full extent of the trial ODD.

For example, for a trial involving a low-speed shuttle in city centre, the scenarios used for testing of the ADS should include pedestrians and other vulnerable road users that the ADS may encounter during the trials.

In order to demonstrate a high level of system safety, Safety Case Creators would need to utilise some form of scenario database to document all the test scenarios used to assess the ADS; this could be generated by the trialling organisation, or could be an external source such as the National Scenario Database (NSDB). Reviewers should satisfy themselves that the database contains no gaps where there are portions of the ODD that are not covered by test scenarios; use of a central repository may make it easier to identify such gaps. In particular, it should be ensured that this coverage of the ODD includes test scenarios that explore the boundaries of the trial ODD.

Depending on the chosen trial route for the trial, the ODD definition may have some special features which are unique to the selected route (e.g. on the day of the trial it might be refuse collection day where the roads are lined by bins). These features would need to be included into the scenario-based testing approach to ensure that the generated scenarios not only capture these unique features, but also that the test programme is able to test the ADS against these unique features. Reviewers should seek evidence in the safety case pertaining to incorporation of such unique features (if any) in the test scenarios generated.

ISO/PAS 21448 (2019) describes methodologies for assuring 'Safety of the Intended Function' (SOTIF) for Advanced Driver Assistance Systems (ADAS) within road vehicle. Whereas functional safety describes methodologies for ensuring the system is suitably robust against hazards caused by system faults, SOTIF is concerned with ensuring that the design of the system is inherently safe when operating without faults. Although not targeted at higher levels of automation, reviewers may find it to be a useful reference, particularly with respect to using testing to uncover scenario permutations that can act as triggering events for hazardous behaviour. It divides scenarios into the following four 'areas':

01
Known safe scenarios

02
Known unsafe scenarios

03
Unknown unsafe scenarios

04
Unknown safe scenarios

The aim of the process is twofold: to uncover previously unknown hazards such that they move from area 3 to area 2, and to perform engineering development to move the area 2 scenarios (known to be unsafe) into area 1 (i.e. to update the vehicle such that the scenarios become safe).

When testing is done according to a scenario database, however, the 'unknown' areas (3 and 4) each have two further subcategories beyond those defined in ISO/PAS 21448 (2019):

a. Untested on vehicle but captured within database

b. Untested on vehicle and not captured in database

The first subcategory contains scenarios that have not yet been discovered to be hazardous for the ADS, but will be discovered to be in due time as the test programme proceeds. However, the second subcategory is more difficult to address, as scenario-based testing will only expose the ADS to scenario permutations that are already known to be possible within that ODD, leaving the possibility of flaws in the ADS remaining uncovered due to gaps in the database. This highlights the importance of a comprehensive database when applying scenario-based testing.

KEY REVIEWER POINT

Until ADS performance has been successfully verified in the full range of scenario permutations applicable to the ODD, operational measures to mitigate risk, such as the presence of a safety driver, should not be removed.

ADVANCED TRIALS

5.3 System safety analysis

Various standards such as ISO 26262 (2018), ISO/TR 4804 (2020) and SAE J3187 (SAE, 2018b) have provided guidance on the application of methods such as FTA, FMEA, HAZOP and STPA as part of the system safety process. Each hazard identification method has advantages and disadvantages, and therefore the selection of an appropriate methodology, with justification, should be documented within the safety case. Depending on the complexity of the trial (including complexity of ODD and level of control in the trial), a combination of these methods may be used.

EXAMPLES AND TEMPLATES

Where a trial involves a safety driver within the vehicle on a closed test track, it may be preferable to adopt a high-level analysis for component level failures (typically an FMEA) with the trained safety driver tasked with responding to any unsafe emergent behaviour. On the other hand, for a trial in an urban city centre with vulnerable road users present, it is important to undertake a detailed hazard identification process and it may therefore be determined to be appropriate to apply more than one analysis method, taking advantage of the differing strengths and weaknesses to reduce the likelihood of hazards not being uncovered or being improperly understood. This may, for example, consist of performing an FMEA combined with either FTA, HAZOP or STPA.

It is important for reviewers to appreciate that there is no standard set of validation targets which the trialling organisations can demonstrate against. In this situation, reviewers should adopt a pragmatic approach to evaluating the test scenarios against the desired behaviour and the defined trial ODD.

KEY REVIEWER POINT

The safety case should include selection of an appropriate method for analysing system safety (e.g. FMEA, FTA), considering the relative merits of the analysis methods and the nature and complexity of the trial.

Reviewers should be satisfied that the trialling organisation has adequately considered what constitutes demonstrating an acceptable level of system safety, bearing in mind that there is currently no industry standard or consensus. The acceptance criteria should be documented as part of this, including:

- The target for how well the scenarios should cover the ODD (i.e. coverage analysis)
- The success criteria for what constitutes a pass in each individual test scenario
- The success criteria for the overall level of performance achieved by the system across all the scenarios.

ADVANCED TRIALS

5.4 Using simulation for evidence generation

Simulation allows a developer the flexibility to test a diverse set of scenarios, especially safety-critical scenarios which may be unsafe for performing upon public roads. Moreover, simulation is relatively efficient with regard to cost and time in comparison to physical testing. However, as per the functional safety processes defined within ISO 26262 (2018), tools which are used for system development and safety evidence generation should be 'qualified', i.e. tested to validate that they perform as intended. Using simulation for testing an ADS therefore requires the developer to ensure that the simulation is representative of the real-world.

While the industry has yet to reach a consensus upon the methods and metrics to be used for validating the accuracy of results generated by simulation software, trialling organisations using simulation-based evidence to support a system safety case should undertake some correlation testing between their simulation environment and real world testing to establish representative behaviour of the ADS.

The rigour of this correlation testing should be dependent on the complexity of trials. For example, a trial with a safety operator on a proving ground may need limited (if any) correlation testing, whereas more complex trials may need to provide evidence for any of the simulation results. Similarly, if the bulk of the evidence of system safety is being generated through physical testing rather than simulation, it may be deemed disproportionate to do an in-depth correlation study for a relatively small proportion of test cases.

Validation of simulation testing should include an analysis of the accuracy in the following areas:

a. Sensor models (the characteristics of the sensor models should match the real sensors, e.g. image distortion from camera lens, noise within radar sensor circuitry)

b. vehicle models (the vehicle dynamics and the actuator responses should be realistic)

c. world/environment models (the surroundings and the weather conditions should be replicated accurately)

Models would typically be evaluated by running identical (or as similar as possible) test cases in simulation and in the real world. This allows an assessment to be made of the correlation of the results; if the results are similar in the two environments, this suggests good correlation. Such correlation test cases would need to be performed in a range of scenario permutations, sampling a range of possibilities within the ODD, as good correlation in one type of test does not guarantee good correlation in another (e.g. realistic behaviour in dry weather does not guarantee realistic behaviour in rain). Models could be validated by testing the entire ADS or by testing an individual element – for example, by injecting a particular control request to the actuators to check that the resulting vehicle path correlates between simulation and physical testing, such that sensor models or ADS path planning are not factors that can affect the results.

KEY REVIEWER POINT

Reviewers should take a pragmatic approach as there is a lack of common understanding on the metrics to be used for validating the simulation platform. Judgement will therefore be needed in order to determine what constitutes an acceptable level of correlation between the simulation and real-world. The safety case should therefore not merely contain simulation test data, but also a safety argument that shows how that data indicates an acceptable level of safety.

5.5 Summary of section 5

For trials that use an on-board safety driver, the analysis of system safety will typically be limited to ensuring the control overrides needed by the safety driver are robust. However, for trials that are dependent upon the system performance being safe, due to the inability of a safety operator to correct mistakes by the ADS, extensive system safety evidence will be needed.

In order to demonstrate that the ADS is able to operate safely, it would be necessary for the test scenarios undertaken to provide coverage of all the possible permutations within the ODD and all the behavioural competencies the vehicle is capable of.

A scenario database can help ensure the test programme provides good coverage of the possible scenario permutations, minimising the likelihood of hazardous system flaws remaining undetected.

Simulation testing can cover a wide range of scenarios efficiently and safely. However, the accuracy of the simulation needs to be validated if the results are to be used as safety evidence.

6.0 The security case

This section examines how to identify security threats, including both physical security and cyber security, how to assess the risk presented by the identified threats, and how to put in place appropriate security controls to manage the risk throughout the trial lifecycle.

6.1 Security introduction

PRELIMINARY TRIALS

6.1.1 Overview of security for self-driving vehicle trials

Security is an important factor that can affect the overall the safety of a self-driving vehicle trial, and therefore the risks that physical or cyber-attacks may pose to the safety of a trial should be considered as an integral part of the safety case. The security lifecycle follows the overall system design and trial lifecycle and, in particular, the safety lifecycle.

In the context of a self-driving vehicle trial, security should be a collaborative responsibility of all stakeholders, including the trialling organisation, testbed, local authorities and any other trial participants; see Section 7.3 for more information on stakeholder engagement. Figure 6.1 shows an overview of a typical trial and the organisational elements involved.

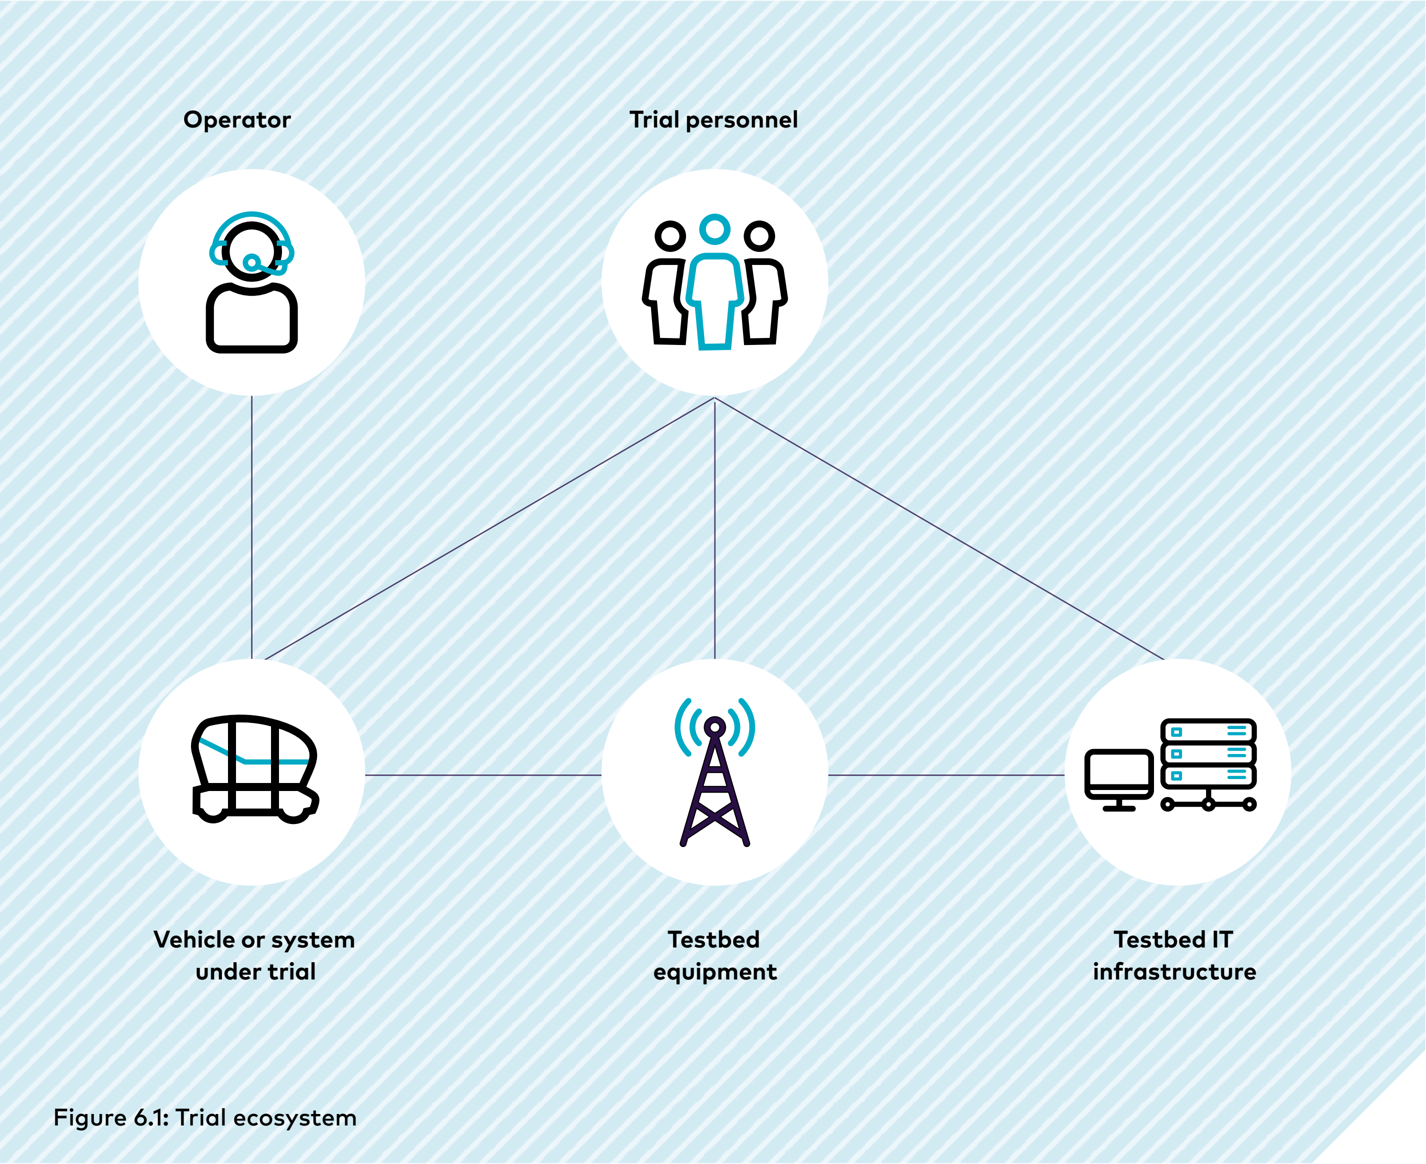

Figure 6.1: Trial ecosystem

Security risks should be managed over the full lifecycle of the self-driving vehicle trial (see Section 7), and cover both the development of the systems being trialled and the operational aspects of the trial whilst it is underway. Responsibilities for security assurance are therefore distributed between all stakeholders involved in the trial, including the manufacturer and suppliers of the systems being trialled, the trialling organisation, the testbeds and any other participants.

The security risks will be different for different types of trial and different testbed environments; for example, proving ground testbeds are likely to be exposed to different threats compared to public road testbeds. Nevertheless, there are a number of security risks and associated best practices that are common between trials, such as ensuring vehicles are locked when not in use and ensuring good cyber hygiene in line with recognised good practice.

Physical security aspects include the trial testbed environment, its associated infrastructure, the route, the trial personnel – whether directly participating in the trial on the testbed or 'behind the scenes' in a remote location – and the public (where applicable).

Cyber security considers all computer-based, electronic and telecommunications systems involved in the trial, including systems associated with the testbed and trial operator, the electronic systems of the vehicle, systems enabling communications, control and monitoring and any remote systems.

Security risk management should also consider the interactions between physical and cyber risks, recognising that a weakness in the physical security of some aspect of a trial can enable a cyber-attack, and vice-versa.

KEY REVIEWER POINT

Safety Case Reviewers should therefore be satisfied that:

- security has been considered within the safety case
- the security case considers the whole development and testing lifecycle of the vehicle
- physical security and cyber security risks have been considered
- appropriate stakeholders have been involved in the process.

PRELIMINARY TRIALS

6.1.2 Proportionality to the trial complexity

As with evidencing operational safety and system safety, the extent of analysis work undertaken, and the evidence documented as part of the security case, should be proportionate to the complexity of the trial. Where there is a safety driver onboard the vehicle and able to take corrective action, the safety case will typically be based upon operational safety; whether errors in the vehicle behaviour are triggered by functional safety (i.e. system faults), safety of the intended function (i.e. inherent limitations in the system design) or cyber security issues, as long as it can be confirmed that the safety driver can correct the error consistently and safely, the error can be tolerated. Therefore, the level of safety evidence expected within the security case would be relatively brief and high level.

The exception to this is where cyber security breaches could compromise the override mechanisms that the safety driver would use to correct the vehicle; these overrides must be robustly assured, including from a security perspective, if the ability of the safety driver to correct ADS mistakes is to be claimed as a component of the safety argument. Furthermore, it must be remembered that the scope of security goes beyond preventing undesired control inputs, and the security case should consider all security risks that could affect the trial.

In contrast, for advanced trials where it is not feasible for human oversight of the ADS to be used as a safety measure (e.g. remote monitoring), far more detailed analysis and evidence would typically be required to ensure that the system is suitably robust against hazardous cyber security breaches such that human oversight is not required. Furthermore, where remote safety operator intervention is claimed as a safety mitigation within the safety argument, the integrity and security of the communications link between the vehicle and operator (in both directions, i.e. vehicle feedback to operator and operator inputs to vehicle) would need robust assurance.

KEY REVIEWER POINT

Reviewers should look for evidence that the trialling organisation has considered proportionality and provided and appropriate level of detail in their security case, bearing in mind the vehicle capabilities, trial ODD and the level of controllability by a safety operator.

6.1.3 Impact of security upon safety

Security should not be seen as an end in itself but as a means to avoid harm or other adverse events due to intentional actions such as malicious attacks. In the context of self-driving vehicle trials, such adverse events can be grouped into the impact categories shown in Table 6.1.

IMPACT CATEGORY	SECURITY OBJECTIVE
Safety	To ensure the safety integrity of the vehicle systems. To minimise the safety risk to vehicle occupants, trial personnel, other testbed users and members of the public.
Privacy	To protect the privacy of vehicle occupants, trial personnel and other testbed users.
Financial	To protect against financial loss due to fraudulent commercial transactions, theft of vehicles, and physical damage to the vehicle, testbed and other assets.
Operational	To maintain the intended operational performance and maintenance of the vehicle and trial environment

Table 6.1: Example security impact rating classifications

KEY REVIEWER POINT

The security case should show evidence of consideration of the four impact categories (safety, privacy, financial, operational), with a focus on the safety category.

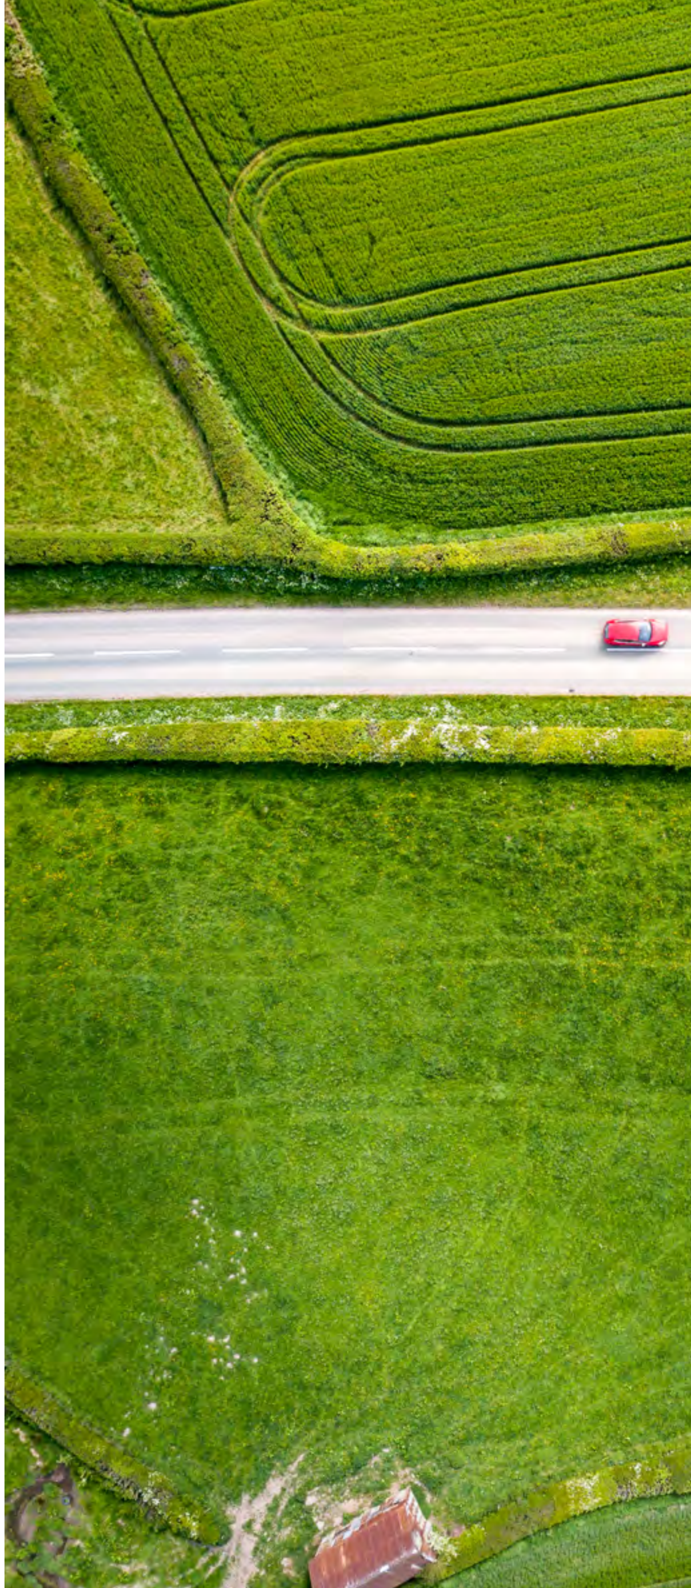

6.1.4 Trial approvals regarding security

Evidence of security risk management of the systems under trial should be built up during their development, with appropriate review and sign-off, before use within the trial. Sign-off is expected to be provided by an appropriate combination of:

- certification of systems under trial against appropriate standards or regulations (for example UNECE Regulation 155 for vehicles (UNECE, 2020)
- local authority sign-off processes
- consideration of other users of the proving ground or testbed, for example through exclusive use agreements.

KEY REVIEWER POINT

There should be an appropriate process for relevant stakeholders to sign-off the security case.

6.1.5 Relevant regulations and standards

The following regulations and standards are applicable to security of self-driving vehicle trials, and may have been used or referenced by the creators of a safety case:

- SAE J3061 Cyber security Guidebook for Cyber-physical Systems (SAE, 2016)
- ISO/SAE DIS 21434 (ISO/SAE, 2020)
- IEC 62443 (2018)
- ISO/IEC 27000 series of standards (ISO/IEC, 2018)
- BSI PAS 1885 (2018)
- BSI PAS 11281 (2018)
- ETSI C-ITS standards (ETSI, 2010, 2012, 2017a)
- UNECE WP.29 vehicle type approval regulations regarding cyber security (UNECE, 2020).

It is acceptable for trialling organisations to 'tailor' their use of a standard, i.e. to deviate from it where appropriate, with justifications recorded.

KEY REVIEWER POINT

There should be evidence of reference to and application of relevant regulations and standards within the security case, although reviewers should not require trialling organisations to follow any one particular standard.

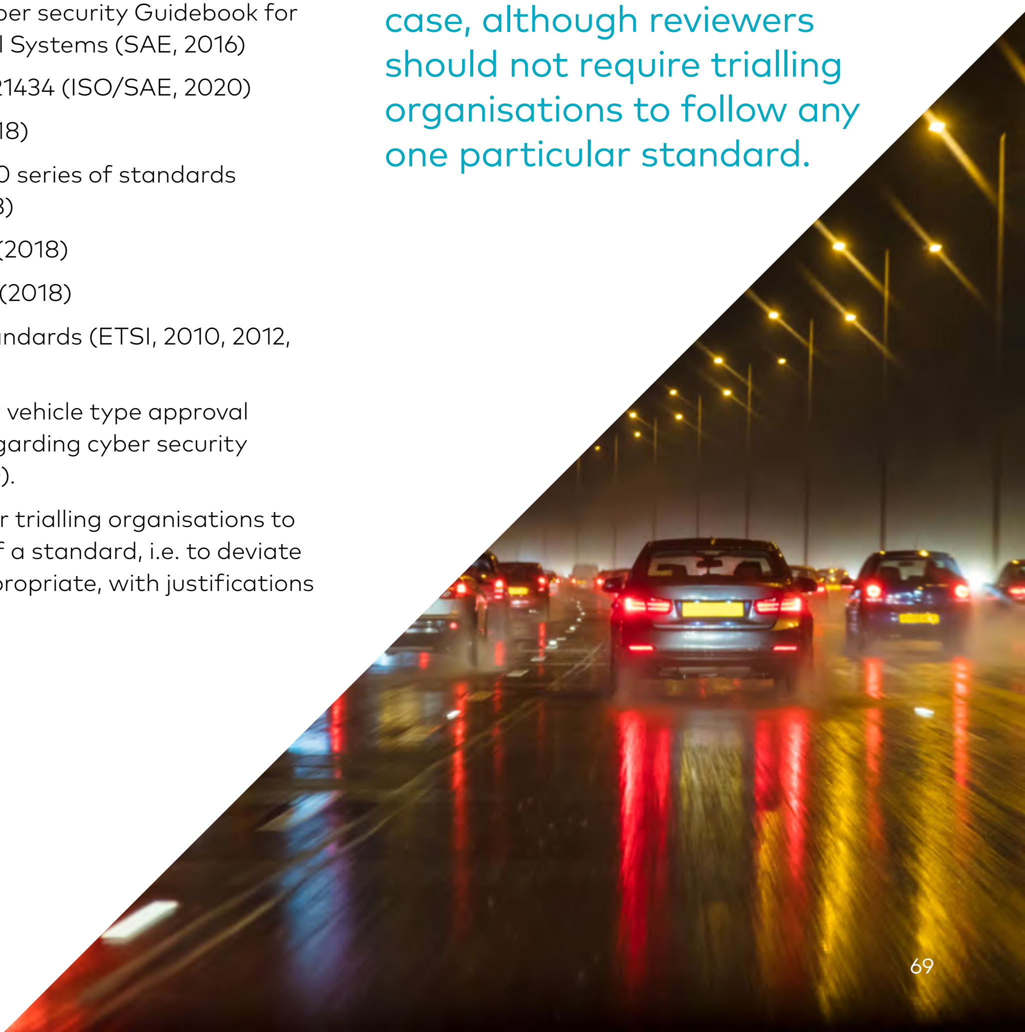

6.2 Management of security

6.2.1 Security management plan

In addition to technical cyber security engineering and operational activities, the organisational management of security is critical to provide the processes and governance structures to support security assurance. In the context of a self-driving vehicle trial, management of security can take the form of a security policy which acknowledges security risk related to the trial ecosystem.

The approach to security for a trial, as defined in the security policy, should be implemented by a more detailed trial security management plan which defines the physical security and cyber security activities for the duration of the trial, to complement the trial safety case. The security management plan can include engineering procedures, methods, design rules and operational guidelines. It may, as appropriate, draw upon or reference more specific process documents of each of the trial participants and describe how security will be managed collectively for the trial.

KEY REVIEWER POINT

The safety case should provide evidence of the existence of a security management plan for the trial, and evidence that this plan has been followed up to the point in time where the safety case is being reviewed.

6.2.2 Information security management

The security management plan should also identify how the participating organisations will identify, classify and manage information during the trial, and can include requirements for information security management.

KEY REVIEWER POINT

The safety case should provide evidence of the approach to information security management taken by the trial and how this will be managed by the participating organisations.

6.2.3 Information sharing and continuous improvement

Since the threat landscape is continuously evolving, it is important that relevant security information is shared between testbeds, trialling organisations and other stakeholders to enable continuous improvement in security processes and measures. Testbeds, trialling organisations and other trial participants should identify circumstances in which information should be shared and with which other parties. For all forms of information sharing, appropriate information security management should be in place to ensure sensitive details of threats or vulnerabilities are handled confidentially.

KEY REVIEWER POINT

Where testbed policy requires sharing of security information, reviewers should ensure that discussions have taken place on this between the testbed and the trialling organisation, and that a suitable plan for data sharing has been documented.

6.3 Security risk management

6.3.1 Risk management approach

Security risk can be defined as the combination of the likelihood of occurrence of a successful attack on a system (be that physical or cyber, or both) and the potential impact on stakeholders.

The dynamic nature of security risk means that the threat landscape is continuously evolving, and new forms of attacks and system vulnerabilities are frequently discovered. Risk management for cyber security should be carried out iteratively during the lifecycle of the trial to assess, evaluate and treat risk. An initial risk assessment should be carried out at the start of development and planning of a self-driving vehicle trial, but risk assessment should also be reassessed, and updated where appropriate, as additional information becomes available. Owners of security risks should also be identified, including from an insurance perspective.

KEY REVIEWER POINT

The latest version of the Security Risk Assessment should be in evidence (this may be the original assessment or may have evolved by the time of review), together with evidence of a process to allow iteration to capture newly identified risks.

The nature of security risk also means that not all the information required to assess risk is available at every point. For example, it is usually possible to identify the end consequences of a threat being realised at a much earlier stage than the attack methods that can be used to realise those threats, since the latter are highly dependent on the system design and implementation.

Security risk management is typically based on the following activities (see also Figure 6.2):

- Define system scope and boundaries (i.e. system architecture)
- Identify assumptions on the system and operational environment
- Identify assets
- Identify and analyse threats and vulnerabilities
- Risk assessment
- Risk treatment and specification of controls/mitigations

A security risk log should be documented to capture the identified assets, threats and vulnerabilities, together with their associated risk assessments and treatments.

KEY REVIEWER POINT

The safety case should contain evidence that the activities of this process have been carried out to an adequate level and that the identified risks are appropriately treated such that residual risks are acceptable and reduced to be ALARP.

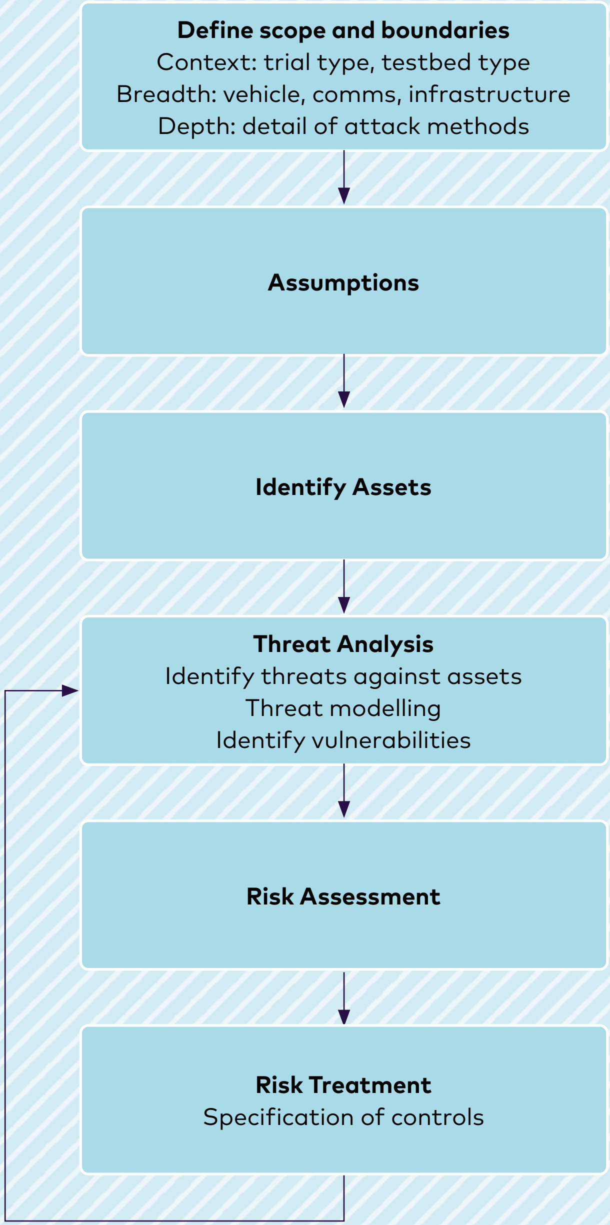
Define scope and boundaries
Context: trial type, testbed type
Breadth: vehicle, comms, infrastructure
Depth: detail of attack methods
Assumptions
Identify Assets
Threat Analysis
Identify threats against assets
Threat modelling
Identify vulnerabilities
Risk Assessment
Risk Treatment
Specification of controls

Figure 6.2: Security risk management process

6.3.2 Scope and boundaries

The scope of the risk assessment should be determined in terms of both

a. the breadth of the analysis in relation to the extent of the trial, and

b. the depth of the analysis with respect to the level of detail in which to consider the threats.

The breadth and depth of analysis will depend upon the type and complexity of the trial, e.g. public road versus controlled test track. For example, a trial on a public road testbed may be more exposed to attack with a wider range of stakeholders potentially at risk, and as such need to consider a wider range of threats and a deeper analysis to identify potential vulnerabilities.

Often the potential safety impacts of a physical or cyber-attack can coincide with those due to hazards identified in the safety risk assessment. Therefore, it is important that there is interaction between the security and safety risk management activities.

KEY REVIEWER POINT

Exclusions, assumptions, dependencies and caveats relating to physical and cyber security should be clearly specified.

6.3.3 Assets

An asset is something of value (actual or perceived) to one or more stakeholders or actors. Therefore, in order to identify the assets in scope for a trial, it is necessary to identify the relevant stakeholders who could potentially experience losses if a threat were to be realised against an asset. In some cases, assets can be elements of the system under trial, trial equipment or elements of the testbed. The relationship between these terms is illustrated in Figure 6.3.

Examples of assets in a self-driving vehicle trial include:

- vehicles
- testbed infrastructure, for example communications units
- data sent to and from vehicles over communications links
- data involved in the functioning of an automated driving system
- personal data of trial participants.

Relevant security properties of the assets can also be identified. For example, the integrity of communications data should be preserved for the entire duration of the trial, even in the case of an incident.

KEY REVIEWER POINT

The safety case should document what assets have been identified as having the potential to be threatened.

6.3.4 Threat identification

Having identified the assets requiring protection by the relevant stakeholders, the threats against those assets can be identified. At this stage the threats are identified regardless of any mitigation measures in place or planned.

A security related threat can be seen as any potential source of damage to the system elements, assets and/or stakeholders, in terms of safety, privacy, financial or operational impacts, that could result from the exploitation of one or more vulnerabilities of a product, process or service by a threat actor in order to achieve a particular attack objective.

A popular method for identifying threats is STRIDE (an acronym of 'spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege'), where consideration is given to the applicability of each category of threat to each element of the system (Shostack, 2014).

Other techniques can be utilised to analyse the identified threats, such as Attack Tree Analysis (ATA), Cause-Consequence Analysis (CCA), influence diagrams. These can be combined with safety analysis techniques such as FTA, FMEA, HAZOP, etc. (see Section 5).

KEY REVIEWER POINT

The safety case should contain evidence that an appropriate method has been used to identify threats.

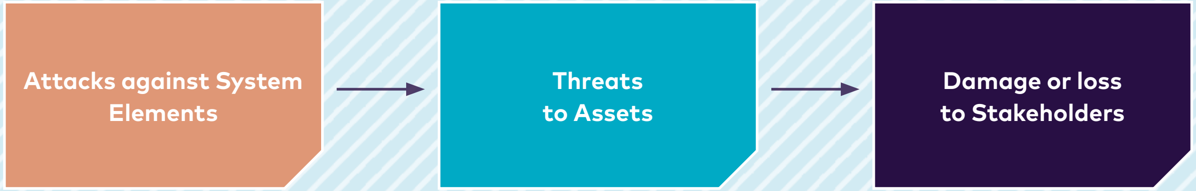
Attacks against System Elements
Threats to Assets
Damage or loss to Stakeholders

Figure 6.3: Relationship between assets, threats, attacks and damage

6.3.5 Risk assessment

The Risk Assessment involves rating the two components of risk for each threat:

- **Impact** – a rating based on the level of impact of the threat upon safety, financial, privacy or operational aspects.
- **Likelihood** – a rating based on how difficult it is to mount a successful attack or exploit a vulnerability. Likelihood is a problematic concept for security risk, due to the inherent non-probabilistic nature of attacks; therefore, a proxy for likelihood such as attack feasibility is often used.

These two ratings are then combined resulting in a risk value. Table 6.2 shows an example security risk matrix where risk is determined as a function of impact and likelihood. Acceptable risk is determined by bands with the same risk value. It is necessary to demonstrate that the security risks have been reduced to an acceptable level.

IMPACT

		NEGLIGIBLE	MODERATE	MAJOR	SEVERE	CRITICAL
LIKELIHOOD	Very high	1	3	4	5	5
	High	1	3	3	4	5
	Medium	1	2	3	3	4
	Low	1	2	2	3	3
	Very low	1	1	1	1	1

Table 6.5: Example security risk matrix

6.3.6 Security controls

The most appropriate security controls for a particular trial will vary according to the threats identified, the type of trial and testbed, the vehicles and vehicle systems involved, and the involvement of personnel and the public in the trial.

For prototype vehicles, some cyber security risks can be effectively managed through the implementation of operational controls as described in Section 4, such as a safety operator with emergency override. However, this may not provide effective mitigation in all cases - for example, some cyber-attacks will not be perceptible by a safety operator (either remote or in the vehicle). In these cases, further security controls may be required to prevent or detect attacks.

Cyber-attacks that could compromise the safety operator's ability to override the vehicle (e.g. by making it more difficult to override the steering or by preventing the system being switched into a standby state) are another example of threats that even a safety driver on board the vehicle would not be able to react safely to, since they affect the robustness of the very mechanism they would need to use in response.

One possible solution would be having an emergency cut-out button that physically breaks an electrical circuit such that the ADS is incapable of providing control inputs to the actuators ; however, this may not be appropriate for all system architectures or trial types.

KEY REVIEWER POINT

The nature of security controls should be clearly stated; for example, whether they are to personnel or operators, physical or cyber aspects, and how they are to be implemented in the system or trial environment. Statements should be in evidence (in the risk log and elsewhere) to demonstrate that the selected security controls are adequate and effective.

6.4 Effectiveness of controls

It is important that the effectiveness of the implemented security controls is verified and validated in order to ensure the risk assessment is as complete and accurate as possible. This can involve independent review of the design and implementation of the vehicle systems and trial infrastructure, as well as security testing.

For cyber security, a range of testing techniques can be applied to verify the effectiveness of the implemented cyber security controls; for example, vulnerability scanning, fuzz testing and penetration testing. The results of these tests can be used to update the risk assessment to reflect whether the actual implementation of the controls sufficiently mitigates the risk. Table 6.3 describes typical cyber security testing methods that can be used to verify the effectiveness of cyber security controls and identify vulnerabilities.

KEY REVIEWER POINT

The evidence that such testing has been carried out should be included as part of the evidence to support the security argument in the safety case.

EXAMPLES AND TEMPLATES

TECHNIQUE	DESCRIPTION
Vulnerability scanning	Testing, usually automated, of a system for instances of known cyber security vulnerabilities. Vulnerability scanning tools exist for many software and network technologies and are an effective way to quickly find known issues, although they are less effective for finding unknown or system-specific issues
Fuzz testing	A method for identifying weaknesses that could potentially be exploited by testing a system with intentionally invalid or malformed input data. The input data can be generated by a combination of random, systematic or adaptive methods, and the effect on the system is monitored to determine any exploitable cases
Penetration testing	A method in which the tester tries to attack the system by adopting similar tools and techniques to a real attacker. This approach is time consuming and is not feasible to apply exhaustively, but it is an effective way of identifying previously unknown vulnerabilities and exploring how they could be exploited

Table 6.3: Summary of typical cyber security testing methods

6.5 Operational monitoring and incident management

Bearing in mind the increasingly rapid pace of technological change and the ingenuity of would-be attackers, new threats will emerge that could not be foreseen during initial design. The trial should therefore implement approaches to detect, understand and respond to incidents that may occur as part of ongoing risk management.

Different trial environments will need different approaches to incident planning and management. The appropriate response actions to a security incident need to be considered as part of the overall incident/ emergency preparedness and planning (as described in Section 7.2). Such incident planning should consider not only the immediate response and remediation activities but also what to do after an incident, appropriate communication activities to inform both the trial stakeholders and the general public, and how lessons learned are fed back to improve future incident responses.

KEY REVIEWER POINT

The security case for the trial should include evidence that an appropriate cyber security incident management process has been planned.

6.6 Assurance of security

6.6.1 Security assurance arguments

A security assurance argument forms an important part of a safety case for a self-driving vehicle trial. Since intentional attacks are a potential cause of safety hazards, the management of security risks needs to be an input to the overall safety risk management and assurance process.

The safety case should contain evidence resulting from the security risk assessment to demonstrate that appropriate and effective security controls have been put in place to mitigate against risks of physical and cyber-attacks to the vehicle, test equipment or trial infrastructure, such that these risks have been managed effectively.

KEY REVIEWER POINT

The security case and security assurance argument for the trial should cover both the argument for 'product security' (related to the vehicle systems being trialled) and 'operational security' (related to the testbeds and trial management).

As was described for the overall safety case (see Section 2.1.2), the security case can use approaches such as GSN (2018) to provide structure to the assurance argument and a clearer view of dependencies and associated evidence. The argument and supporting evidence may also highlight any residual risks and areas that require further evidence to be provided that may be generated in later phases of the trial.

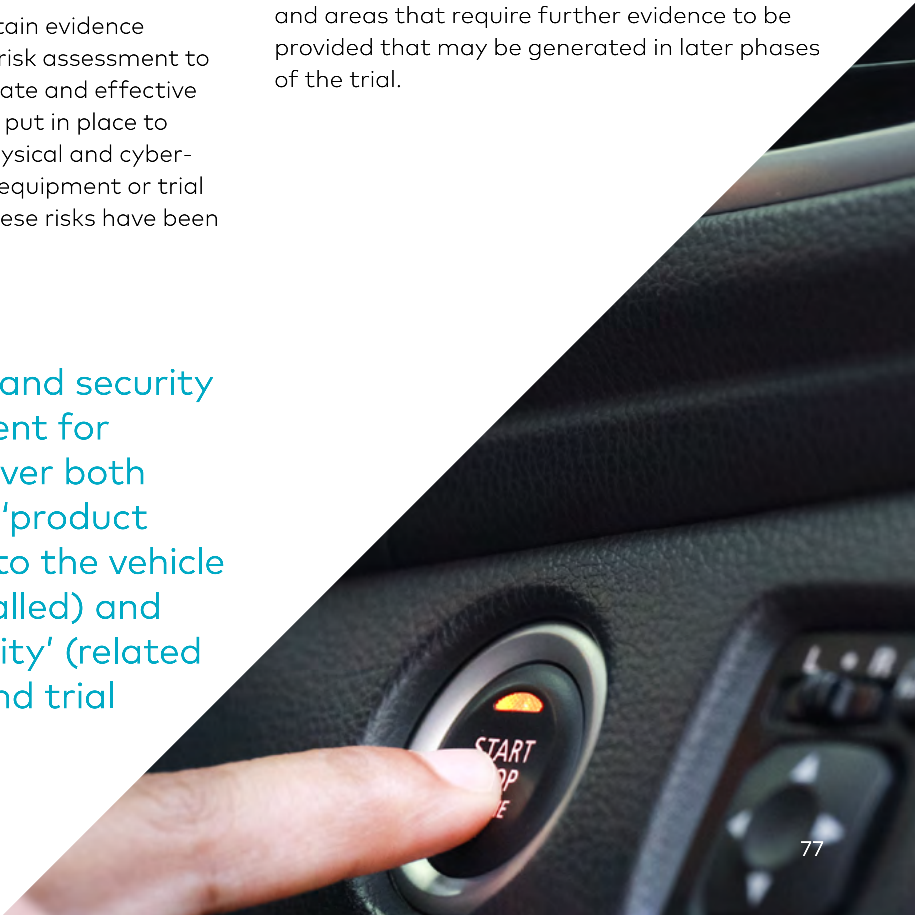

6.6.2 Progressive security assurance

The assurance of both safety and security should be progressive throughout the project, including during trials phase, and aligned with key project stage gates. Due to the influence of adaptive human adversaries, security related threats evolve continuously, and as such there also needs to be a process to regularly review and update the security arguments and supporting evidence.

KEY REVIEWER POINT

The safety case should make reference to how the security argument it contains will be updated over time. This should be aligned to the project plan and the safety case documentation for system and operational safety.

6.7 Summary of section 6

The safety case should consider security threats, including both physical security and cyber security.

The trial assets should be defined, and then a threat analysis should identify and document the different threats to these assets. Methods such as STRIDE can help to identify threats.

A risk assessment should be performed to assess the relative risks posed by the different threats. This will then be used to prioritise security controls for the trial.

An ongoing process for managing security risks should be in place to ensure that new or changed risks are identified and that controls are updated accordingly.

Bolventor
St Neot
St Cleer
Siblyback
Lake

7.0 Process considerations

This section examines how the safety case should include evidence of suitable processes to allow safety to be managed on an ongoing basis as the trial progresses, such as a means to report incidents and a means to trigger safety case updates. Processes are also examined for consultation with stakeholders and for ensuring compliance with documents such as regulations and standards.

7.1 Change control

7.1.1 Purpose of a change control process

A trialling organisation may need to make changes to systems or operational procedures. These changes may be necessary, such as when a fault is identified, or they may be desirable, such as when working towards more advanced systems or operation in less controlled environments.

It is important that changes do not increase the level of risk, or invalidate any evidence or previously held assumptions that risk is ALARP within the operational safety case. A change control process, in accordance with BSI PAS 1881 (2020), should aim to ensure no changes to a system, process or activity are made without consideration of potential risk, and should also ensure that any changes made are documented within the safety case. Logging all changes made through the change control process will also create an audit trail.

7.1.2 Key elements of a change control process

Changes that may trigger the change control process can arise from systems (hardware, software and data) or operational procedures (processes, activities, roles and responsibilities), wherever these changes could affect the safety of the trial.

Changes may be proactive, for example:

- as the next stage of a pre-determined development plan
- due to the identification of a fault or gap during a periodic review
- in response to requests from a stakeholder

or reactive, for example:

- as a result of an incident
- due to the failure of a test
- due to a change in a regulation, or standard.

The change control process should consist of a number of key elements:

- identification of change
- initial determination of the impact to safety
- conducting a full impact assessment
- development of an implementation and monitoring plan
- creation of an approvals process
- updating of any safety case documentation (if required)
- implementation, monitoring and review of change.

The process should also identify the key roles and responsibilities that are needed to ensure this process is correctly followed.

These key roles might include:

- **Change Proposer** – details and justifies the reason for the change. This can be anyone involved in the trials; for example, safety managers, engineers or external stakeholders. Everyone should feel comfortable in raising concerns and suggesting possible improvements.
- **Change Owner** – trial team member tasked with taking the change through the change control process, coordinating the preparation of an impact assessment, creating implementation/monitoring plans, updating safety case documentation and seeking approvals. Examples of change owners may include software engineers, safety engineers and technical leads.
- **Approver** – trial or safety manager responsible for checking that processes are being correctly followed and for signing off updates made to the safety case documentation by the Change Owner.

An example of a change control workflow, which includes the key elements of a change control process, is presented in Figure 7.1.

KEY REVIEWER POINT

Safety Case Reviewers should be satisfied that the trial has a process in place that:

- provides a feedback mechanism so lessons can be learnt
- allows changes to the proposed trial and/or system to be assessed and documented
- assigns clear roles and responsibilities.

EXAMPLES AND TEMPLATES

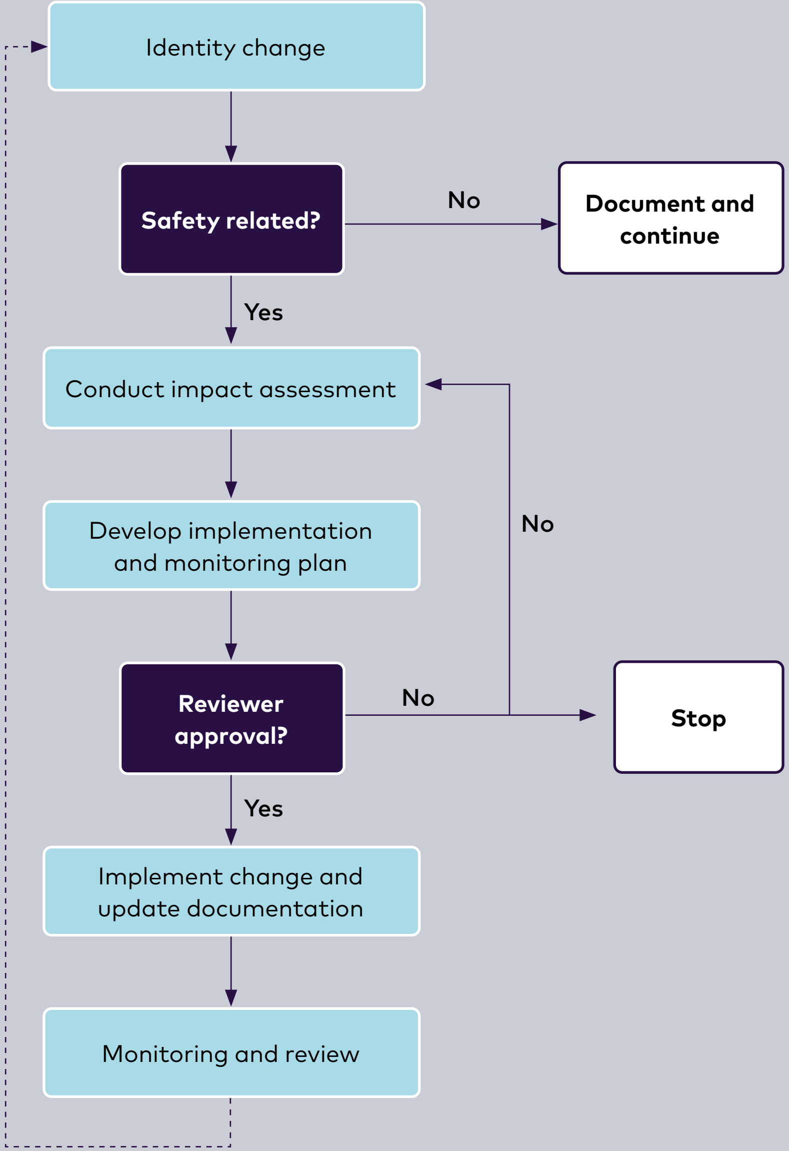
Identity change
Safety related?
No
Document and continue
Yes
Conduct impact assessment
Develop implementation and monitoring plan
No
Reviewer approval?
No
Stop
Yes
Implement change and update documentation
Monitoring and review

Figure 7.1: Example change control process with key elements

7.2 Continuous improvement: monitoring and reporting

7.2.1 The purpose of a continuous improvement process

A trialling organisation should have in place a system for continuous improvement, through a cyclical process of monitoring and reporting, that can gather the evidence to inform more accurate assessments of risk posed, thereby validating decisions that have been made. This evidence could be qualitative or quantitative, although it is unlikely that there will be sufficient data to form statistically significant conclusions within limited scale research projects. Collection of such data could subsequently support the safety argument for deploying the system in more complex trials.

The frequency of this cycle of continuous improvement should reflect the complexity of the trials, with consideration of the trial duration. For example, testing due to be completed within a single day may benefit from a trial team safety briefing before tests commence and another at an appropriate mid-point, whereas trials that continue across a number of weeks may initially benefit from similarly frequent monitoring and additional feedback opportunities, but feature a progressive reduction in frequency as the trial progresses and confidence in operational procedures increases.

7.2.1.1 Monitoring

A monitoring plan should demonstrate how a trialling organisation plans to observe trials and tests and collect data from vehicle systems, complementary infrastructure and trial staff to:

- gather evidence to validate assumptions and risk decisions
- assess the effectiveness of controls and risk mitigation measures
- understand ADS safety performance
- ensure safety is maintained
- identify near misses
- assist with incident investigation
- learn about interactions with other road users.

KEY REVIEWER POINT

Reviewers should be satisfied that there is an adequate plan in place to monitor the safety of the trials as they progress.

7.2.1.2 Incident investigation and reporting

An 'incident' is an unintended and undesired event such as a collision, a near miss, a technical malfunction with potential safety implications, a transgression of road traffic laws or a security breach. An incident investigation and reporting plan should exist to document such events, and should include both processes that are proactive (routine reviews according to a planned cadence) and processes that are reactive (as a result of a more serious incident). This is to ensure that the circumstances surrounding an incident are investigated, such that measures can be implemented to minimise the potential for reoccurrence. Lessons learned and changes made should then be communicated to relevant stakeholders.

The process will likely involve different process paths depending on appropriate incident classification (see example classification in Table 7.1), but in each case should demonstrate how a trialling organisation plans to observe trials and tests and collect data to:

- analyse safety related data
- identify areas of improvement
- investigate incidents
- communicate lessons learned
- inform stakeholders
- fulfil legal requirements such as RIDDOR (HSE, 2013).

EXAMPLES AND TEMPLATES

INCIDENT LEVEL	INVESTIGATION TYPE	EXAMPLE
01 Moderate	Periodic review	Fault code triggered
	Internal investigation – pause trial	Mechanical breakdown Fault code triggered repeatedly Hand back of control t o safety driver
02 Substantial	Internal investigation – stop trial Landowner/operator investigation	Near miss where the hazard had the potential to cause harm
03 Severe	Internal investigation – stop trial Criminal investigation Civil investigation (e.g. insurance) Research investigation Landowner/operator investigation	Collision

Table 7.1: Example incident classifications

KEY REVIEWER POINT

The safety case should provide reviewers with evidence that a suitable process to classify and investigate incidents is in place. There is no obligation for trials to use the example in Table 7.1, provided that a suitable equivalent is documented.

The process defined to investigate and respond to incidents may include the use of a safety steering group to oversee the process and make key decisions on responses to incidents. The number of persons included within the safety steering group, and the job roles and specialisms represented, should be selected according to the complexity and needs of each trial.

Examples of appropriate roles include:

- trial manager
- trial safety manager
- lead engineer developing system
- senior manager from within the trialling organisation
- local authority, road authority or testbed representative
- external consultant (where specialist and/or impartial input is required).

A trialling organisation should have an incident reporting template that allows any trial team member (without significant understanding of the incident) to collect the necessary information that can enable an investigation of an incident, after the immediate incident response.

This information should include the following:

- Trial team member completing report – name, job title, organisation, role within trial team, contact information
- Trial Manager and Safety Manager at time of incident – name, job title, organisation, contact information
- Safety Operator and Trial Engineer at time of incident – name, job title, organisation, contact information
- Other vehicle occupants – name, job title, organisation, contact information
- Third parties involved – name, contact information, insurance details
- Date and time of incident
- Location of incident – testbed details, road names, GPS coordinates
- Incident classification in accordance with predefined incident reporting plan (based on understanding at the time)
- Reasons for incident (based on understanding at the time)
- Operating mode of vehicle at time of incident (based on understanding at the time) – for example manual driving, automated driving, remotely operated
- Actions taken
- Injuries and damage details and any RIDDOR reporting if required
- Emergency services attended – details of which services attended, contact information
- Whether vehicle recovery was required
- Insurers informed – data communicated
- Data collected, stored and communicated

7.3 Stakeholder consultation and engagement

7.3.1 Purpose of the stakeholder consultation plan

The purpose of a stakeholder consultation plan is to:

- **inform** of upcoming trials
- **educate** about the technology and the controls that are in place to ensure safety
- **collect** information and local knowledge
- get **permissions** to operate trials and tests
- get appropriate **derogations** for trials and tests
- get any required **licencing** to operate
- **understand** and put in place stakeholder requirements.

Public consultation will be of particular importance where members of the public participate in the trial (e.g. as passengers). Such trials should consider whether the way the public are involved is appropriate and ethically acceptable, e.g. via an ethics committee reviewing the planned activities. Consultation with members of the public can help inform and validate such decisions.

7.3.2 Key elements of the stakeholder consultation plan

As part of producing a stakeholder consultation and engagement plan, a list of all stakeholders should be compiled, including:

- the nature of their involvement
- how and when they will be contacted
- the methods of engagement
- the information shared or received.

Table 7.2 presents examples of possible stakeholder types, and the possible reasons for consultation or engagement. Note that not all stakeholders will be applicable to all trials; for example, some will be exclusively applicable to proving grounds or to public roads.

KEY REVIEWER POINT

Safety Case Reviewers should be satisfied that the key stakeholders for the trial have been identified, and contacted where appropriate. The reviewer would not need to be provided with evidence of any form of 'sign-off' from these stakeholders, unless this is specifically required as part of a formal process to access a testbed.

STAKEHOLDER TYPE		PURPOSE
Trial Team	Marshal/steward Safety manager Software engineer/Automated Control System Operator Safety driver Trial manager Researchers	Information, training, requirements, incident response, compliance
Landowners and operators	Testbed/Proving ground Track control	Permissions, requirements
	Local Authorities: County Councils, District Councils, Unitary Authorities, Metropolitan Districts Transport Authorities Road/Highway Authorities	Information, permissions, derogations, licencing, compliance, requirements
Licencing and regulators	Traffic Commissioners for Great Britain DVLA DVSA VCA Transport Authorities	Information, permissions, licencing, compliance, requirements
Emergency services	Incident response team Track Control Vehicle recovery Research ethics panel Police Ambulance Service Fire brigade Traffic Officers	Information, incidence response, requirements
Trial participation	Trial participants Research ethics panel	Information, permissions
Local services	Local transport operators Local services/amenities Local construction sites Local businesses	
Local residents	Local residents and other members of the public Local road users	Information, local knowledge
Funders	Consortium members InnovateUK	Information, requirements, incident response
Other stakeholders	Insurers Safety Advisory Groups	Information, requirements, incident response

7.3.3 Guidance on reviewing a publicly available safety case for public road trials

A publicly available safety case is an abridged version of the overall safety case, optimised for public consumption. This safety case should be a simplified, single document, written in plain English without unnecessary or unexplained technical language. The aim of this safety case document is to educate the public on the trial activity and provide reassurance that reasonable steps have been taken to assure trial safety. The Code of Practice for Trialling Autonomous Vehicles (CCAV, 2019b) and BSI PAS 1881 (2020) both recommend that a publicly available safety case should be made available before conducting trials in public domains.

EXAMPLES AND TEMPLATES

Sufficient details should be included to allow members of the public to get a broad understanding of the trial activities. Content of this safety case document may include:

- an overview of the trial activity
- reassurance that trial activity can be performed and managed safely
- a high-level overview of the vehicle and automated driving system
- information about the trial test area
- the roles and responsibilities of the trials team
- details about trial compliance with regulations, standards etc.
- information about trial milestones
- stakeholder communication and engagement plans
- points of contact.

KEY REVIEWER POINT

Safety Case Reviewers may wish to review a copy of any publicly available safety case produced by the trialling organisation. However, it is not essential to the safety of the trial that the public safety case is reviewed, and therefore reviewers should consider whether this level of oversight is proportionate. Safety Case Reviewers who are local to the test area may be able to assist trialling organisations in communicating the publicly available safety case to local businesses and residents as appropriate.

7.4 Compliance

Trialling organisations should demonstrate that they will conduct any trials in accordance with relevant UK law, or have received appropriate permissions or derogations for any cases of non-compliance (e.g. the use of bus lanes). For operations within a proving ground environment, UK road traffic and vehicle regulations may not be applicable, but the safety impacts of non-compliance with relevant data and cyber security law should be considered, addressed and detailed within the operational safety case. In all cases, the trial safety case should show sufficient management of safety to be in compliance with all UK Health and Safety legislation.

For each statement of compliance with a clause, article, regulation, standard or guidance document, the safety case should include the argument and appropriate evidence that can support it. In the case of any non-compliance, details of the necessary permissions or derogations to conduct the trial must be included, in addition to the presentation of the argument and evidence that the risk associated with the operations remain ALARP.

Table 7.3 presents a number of areas of regulation and guidance that should be considered within the operational safety case. This list is not exhaustive, and not all examples will be applicable to all trials; the scope definition within each regulation, code or standard will help clarify applicability. If Safety Case Reviewers require compliance with any particular document in order to access a testbed, this should be made clear to trialling organisations at the earliest possible opportunity. Consideration should be given to providing supplementary guidance to support Safety Case Creators in achieving compliance with any safety processes that are specific to the testbed (or the organisations managing the testbed).

EXAMPLES AND TEMPLATES

AREA	EXAMPLES
Road traffic law	Highway Code The Road Vehicles (Construction and Use) Regulations 1986 Road Traffic Act 1988
Road vehicle in-use regulation	MOT Test Road Vehicles (Construction and Use) Regulations
Cyber security standards	As detailed in Section 6
Incident investigation standards	PAS 1882 (upcoming)
Relevant landowner/operation requirements	GG104 (Highways England's safety process)
Relevant CAM trial standards	Code of Practice for Automated Vehicle Trialling PAS 1881
Health and safety legislation	The Health and Safety at Work etc. Act 1974 The Management of Health and Safety Regulations 1999

Table 7.3: Area of compliance, with examples

7.5 Summary of section 7

Trials should have a change control process that empowers all relevant stakeholders to initiate change and involves all relevant stakeholders in making suitable updates to the safety case and operating procedures.

There should be a plan for how to monitor safety as the trial proceeds, including a means for incidents to be reported, logged and learnt from.

Trialling organisations should consult with affected stakeholders. For public road trials, this should include production of an abridged version of the safety case that is made publicly available.

The safety case should document how relevant documents such as regulations and standards have been reviewed to ensure the trial is compliant with legal obligations and with best practice.

8.0 References

BSI (2020) *CAV Vocabulary v3.0*, available at: https://www.bsigroup.com/en-GB/CAV/cav-vocabulary/

BSI PAS 1881 (2020) *PAS 1881:2020 Assuring the safety of automated vehicle trials and testing – Specification*, available at: https://www.bsigroup.com/en-GB/CAV/pas-1881/

BSI PAS 1883 (2020) *PAS 1883:2020 Operational Design Domain (ODD) – taxonomy for an automated driving system – Specification*. Available at: https://www.bsigroup.com/en-GB/CAV/pas-1883/

BSI PAS 1885 (2018) *PAS 1885:2018, The fundamental principles of automotive cybersecurity. Specification*. Available at: https://shop.bsigroup.com/ProductDetail?pid=000000000030365446

BSI PAS 11281 (2018) *PAS 11281:2018 Connected automotive ecosystems. Impact of security on safety. Code of practice.* Available at: https://shop.bsigroup.com/ProductDetail?pid=000000000030365540

CCAV (2019a) *New system to ensure safety of self-driving vehicles ahead of their sale*, accessible at: https://www.gov.uk/government/news/new-system-to-ensure-safety-of-self-driving-vehicles-ahead-of-their-sale

CCAV (2019b) Code of Practice: Automated vehicle trialling, Centre for Connected and Autonomous Vehicles, Published 6th February 2019. Available at: https://www.gov.uk/government/publications/trialling-automated-vehicle-technologies-in-public

ETSI (2010) *TS 102 731 Technical Specification Intelligent Transport Systems (ITS); Security; Security Services and Architecture*, V1.1.1, September 2010. Available at: https://www.etsi.org/deliver/etsi_ts/102700_102799/102731/01.01.01_60/ts_102731v010101p.pdf

ETSI (2012) *TS 102 940 Intelligent Transport Systems (ITS); Security; ITS communications security architecture and security management*, V1.1.1, June 2012. Available at: https://www.etsi.org/deliver/etsi_ts/102900_102999/102940/01.01.01_60/ts_102940v010101p.pdf

ETSI (2017a) *TS 103 097 Intelligent Transport Systems (ITS); Security; Security header and certificate formats*, V1.3.1, October 2017. Available at: https://www.etsi.org/deliver/etsi_ts/103000_103099/103097/01.03.01_60/ts_103097v010301p.pdf

GSN (2018) Goal Structuring Notation Community Standard Version 2. Available at: https://www.goalstructuringnotation.info/

Highway Code (2019) *The Highway Code*, Department for Transport, last updated 2019. Available at: https://www.gov.uk/guidance/the-highway-code

HSE (2013) RIDDOR - *Reporting of Injuries, Diseases and Dangerous Occurrences Regulations 2013*, Health and Safety Executive. Available at: https://www.hse.gov.uk/riddor/

IEC 62443 (2018) *Security for industrial automation and control systems*. Available at: https://webstore.iec.ch/publication/33615

ISO 26262 (2018) *Road Vehicles – Functional Safety*, available at: https://www.iso.org/standard/68383.html

ISO/IEC (2018) *ISO/IEC 27000 Series, Information technology — Security techniques — Information security management systems*. Available at: https://www.iso.org/standard/73906.html

ISO/PAS 21448 (2019) *Road Vehicles – Safety of the Intended Functionality*, available at: https://www.iso.org/standard/70939.html

ISO/SAE (2020). *ISO/SAE DIS 21434, Road vehicles — Cybersecurity engineering*". Public draft version available at: https://www.iso.org/standard/70918.html

ISO/TR 4804 (2020) *Road vehicles — Safety and cybersecurity for automated driving systems — Design, verification and validation*. Available at: https://www.iso.org/standard/80363.html

RAND (2020) *Safe Enough - Approaches to Assessing Acceptable Safety for Automated Vehicles*, RAND Corporation. Available at: https://www.rand.org/pubs/research_reports/RRA569-1.html

SAE (2016) *J3061_201601, Cybersecurity Guidebook for Cyber-Physical Automotive Systems*, SAE Vehicle Electrical System Security Committee

SAE (2018a) *J3016_201806 Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. Available at: https://www.sae.org/standards/content/j3016_201806/

SAE (2018b) J3187 *Applying System Theoretic Process Analysis (STPA) to Automotive Applications*, available at: https://www.sae.org/standards/content/j3187/

SaFAD (2019) *Safety First for Automated Driving*. Available at: https://www.daimler.com/innovation/case/autonomous/safety-first-for-automated-driving-2.html

Shostack, A. (2014). *Threat Modelling: Designing for Security*, John Wiley & Sons Inc: Indiana.

UL 4600 (2020) *Standard for Evaluation of Autonomous Products*, Underwriters Laboratories. Available at: https://www.shopulstandards.com/ProductDetail.aspx?productid=UL4600

UNECE (2020) *Proposal for a new UN Regulation on uniform provisions concerning the approval of vehicles with regards to cyber security and cyber security management system*. Available at: https://unece.org/fileadmin/DAM/trans/doc/2020/wp29grva/ECE-TRANS-WP29-2020-079-Revised.pdf

ACCELERATING THE SELF-DRIVING REVOLUTION

Zenzic-UK Ltd (i411542 02/21)

zenzic.io