



UL 4600

STANDARD FOR SAFETY

Evaluation of Autonomous Products

This document is only permissible for use by Anonymous User of University of Toronto. 4/16/2020 6:21:36 PM

UL Standard for Safety for Evaluation of Autonomous Products, UL 4600

First Edition, Dated April 1, 2020

Summary of Topics

This is the First edition of ANSI/UL 4600, Standard for Safety for Evaluation of Autonomous Products.

The new requirements are substantially in accordance with Proposal(s) on this subject dated December 13, 2019 and February 14, 2020.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form by any means, electronic, mechanical photocopying, recording, or otherwise without prior permission of UL.

UL provides this Standard "as is" without warranty of any kind, either expressed or implied, including but not limited to, the implied warranties of merchantability or fitness for any purpose.

In no event will UL be liable for any special, incidental, consequential, indirect or similar damages, including loss of profits, lost savings, loss of data, or any other damages arising out of the use of or the inability to use this Standard, even if UL or an authorized UL representative has been advised of the possibility of such damage. In no event shall UL's liability for any damage ever exceed the price paid for this Standard, regardless of the form of the claim.

Users of the electronic versions of UL's Standards for Safety agree to defend, indemnify, and hold UL harmless from and against any loss, expense, liability, damage, claim, or judgment (including reasonable attorney's fees) resulting from any error or deviation introduced while purchaser is storing an electronic Standard on the purchaser's computer system.

This document is only permissible for use by Anonymous User of University of Toronto. 4/16/2020 6:21:36 PM

No Text on This Page

This document is only permissible for use by Anonymous User of University of Toronto. 4/16/2020 6:21:36 PM

UL COPYRIGHTED MATERIAL –
NOT AUTHORIZED FOR FURTHER REPRODUCTION OR
DISTRIBUTION WITHOUT PERMISSION FROM UL



Documents

Document Sharing Hub

APRIL 1, 2020



ANSI/UL 4600-2020

1

UL 4600

Standard for Evaluation of Autonomous Products

First Edition

April 1, 2020

This ANSI/UL Standard for Safety consists of the First Edition.

The most recent designation of ANSI/UL 4600 as an American National Standard (ANSI) occurred on April 1, 2020. ANSI approval for a standard does not include the Cover Page, Transmittal Pages, and Title Page.

Comments or proposals for revisions on any part of the Standard may be submitted to UL at any time. Proposals should be submitted via a Proposal Request in UL's On-Line Collaborative Standards Development System (CSDS) at <https://csds.ul.com>.

UL's Standards for Safety are copyrighted by UL. Neither a printed nor electronic copy of a Standard should be altered in any way. All of UL's Standards and all copyrights, ownerships, and rights regarding those Standards shall remain the sole and exclusive property of UL.

COPYRIGHT © 2020 UNDERWRITERS LABORATORIES INC.

This document is only permissible for use by Anonymous User of University of Toronto. 4/16/2020 6:21:36 PM

UL COPYRIGHTED MATERIAL –
NOT AUTHORIZED FOR FURTHER REPRODUCTION OR
DISTRIBUTION WITHOUT PERMISSION FROM UL



Documents

Document Sharing Hub

No Text on This Page

This document is only permissible for use by Anonymous User of University of Toronto. 4/16/2020 6:21:36 PM

UL COPYRIGHTED MATERIAL –
NOT AUTHORIZED FOR FURTHER REPRODUCTION OR
DISTRIBUTION WITHOUT PERMISSION FROM UL



Documents

Document Sharing Hub

CONTENTS

1	Preface (Informative).....	5
1.1	Goal.....	5
1.2	Scope	5
1.3	Use of this standard with other standards.....	5
2	Scope	6
2.1	Scope summary.....	6
2.2	Elements in scope.....	7
2.3	Scope limitations.....	8
3	Referenced Publications.....	11
3.1	Normative references.....	11
3.2	Informative references.....	11
4	Terms, Definitions, and Document Usage	12
4.1	How to interpret normative elements (Normative)	12
4.2	Terms and definitions (Normative)	16
4.3	Abbreviations and Acronyms (Informative).....	21
5	Safety Case and Arguments	22
5.1	General.....	22
5.2	Safety case style and format.....	24
5.3	Claim and argument sufficiency	28
5.4	Evidence sufficiency.....	32
5.5	Accepted risks	35
5.6	Safety culture	36
5.7	Item scope	38
6	Risk Assessment	40
6.1	General.....	40
6.2	Fault model	41
6.3	Hazards	57
6.4	Risk evaluation	59
6.5	Risk mitigation and evaluation of mitigation effectiveness.....	64
7	Interaction with Humans and Road Users	69
7.1	Human interaction.....	69
7.2	Human communication.....	70
7.3	Interactions with humans and animals	74
7.4	Human contribution to operational safety	82
7.5	Vulnerable road user interaction.....	86
7.6	Other vehicle interaction.....	89
7.7	Mode changes that invoke human safety responsibility	92
8	Autonomy Functions and Support.....	94
8.1	General autonomy pipeline.....	94
8.2	Operational Design Domain (ODD)	96
8.3	Sensing.....	101
8.4	Perception.....	110
8.5	Machine learning and “AI” techniques.....	113
8.6	Planning.....	120
8.7	Prediction.....	123
8.8	Item trajectory and system control	124
8.9	Actuation.....	128
8.10	Timing.....	129
9	Software and System Engineering Processes	130
9.1	Development process rigor.....	130
9.2	Software quality.....	135
9.3	Defect data.....	138
9.4	Development process quality.....	138

10	Dependability	140
10.1	General	140
10.2	Degraded operations	140
10.3	Redundancy	146
10.4	Fault detection and mitigation	152
10.5	Item robustness	157
10.6	Incident response	159
10.7	System timing	170
10.8	Cybersecurity	172
11	Data and Networking	176
11.1	General	176
11.2	Data communications and networks	176
11.3	Data storage	183
11.4	Infrastructure support	186
12	Verification, Validation, and Test	189
12.1	Verification, Validation (V&V), and test approaches	189
12.2	V&V methods	190
12.3	V&V coverage	194
12.4	Testing	197
12.5	Run-time monitoring	204
12.6	Safety case updates	208
13	Tool Qualification, COTS, and Legacy Components	211
13.1	General	211
13.2	Tool identification	212
13.3	Tool risk mitigation	214
13.4	COTS and legacy risk mitigation	219
14	Lifecycle Concerns	221
14.1	General	221
14.2	Requirements/design validation	222
14.3	Handoff from design to manufacturing	224
14.4	Manufacturing and item deployment	228
14.5	Supply chain	229
14.6	Field modifications and updates	231
14.7	Operation	235
14.8	Retirement and disposal	238
15	Maintenance	240
15.1	Maintenance and inspection	240
15.2	Required maintenance and inspections	241
15.3	Non-operational safety	245
16	Metrics and Safety Performance Indicators (SPIs)	246
16.1	General	246
16.2	Metric definition	246
16.3	Metric analysis and response	254
17	Assessment	257
17.1	Conformance assessment	257
17.2	Conformance assessment package	259
17.3	Independent assessment	262
17.4	Conformance monitoring	269
17.5	Prompt element feedback	272

Annex A (Informative) – Use with ISO 26262 and ISO/PAS 21448

A1	Compatibility	274
A2	Safety Case	274
A3	Clause Mapping to ISO 26262:2018	274
A4	Clause Mapping to ISO/PAS 21448:2019	276

This document is only permissible for use by Anonymous User of University of Toronto. 4/16/2020 6:21:36 PM

1 Preface (Informative)

1.1 Goal

1.1.1 This standard is intended to help ensure that an acceptably thorough consideration of safety for an autonomous product has been performed during the design process and will continue to be done throughout the system lifecycle. It does so by emphasizing repeatable assessment of the thoroughness of a safety case.

1.1.2 Conformance with this standard is not a guarantee of a safe automated vehicle. However, conformance with this standard promotes more rigorous engineering in support of a safe automated vehicle. It is also recognized that a safety case is just one of many important parts to a complete safety assurance framework for automated vehicles, and it is expected that this standard will be used in conjunction with other standards and test methodologies defined by standards organizations and regulators.

1.2 Scope

1.2.1 The scope of this standard is a generalized autonomous system standard framework using light autonomous road vehicles as a concrete example. To that end, this version of the standard includes extensive prompt lists applicable to light autonomous road vehicles (both passenger and cargo vehicles). Many of the prompts will apply to other autonomous ground vehicles and even other types of autonomous systems, but no specific attempt has been made to include extensive prompts for other applications, nor to segregate road vehicle prompts from more general prompts.

1.2.2 The approach taken in this standard (UL 4600) is to require a claim-based safety case that encompasses essentially the entirety of the material necessary for safety assurance. The safety case includes a structured set of claims, argument, and evidence supporting the proposition that an item (a vehicle plus all other support contributing to safety) is acceptably safe for deployment. In support of that goal, UL 4600 assessments emphasize ensuring that the safety case is reasonably complete and well formed. In particular, UL 4600 provides guidance to improve consistency and completeness of the safety case. To this end, some best-practice process activities and granular work products are specifically required (e.g., creation of a hazard log). However, no specific overall design process is mandated, nor are there mandates for specific methods used to create the majority of work products (e.g., a V-style development process is not required; any reasonable approach used to create a list of hazards can be acceptable).

1.2.3 This standard does NOT define a process, but rather puts forth assessment criteria to determine the acceptability of a safety case. As such, the ordering of sections, clauses, and prompt elements does NOT imply temporal ordering or other process path dependencies.

1.3 Use of this standard with other standards

1.3.1 This standard is intended to work with existing standards to provide the additional elements necessary to assure that safety aspects of fully autonomous item operation have been considered in a comprehensive manner when creating a safety case.

1.3.2 To the maximum extent practicable, it is intended that developers can take advantage of effort expended and assessment credit gained for conformance to other existing standards. Developers may incorporate materials into their safety case generated as a result of executing processes and generating work products required by other standards.

1.3.3 It is the intent of this standard to be compatible with existing relevant safety standards to the maximum extent practicable, and in particular avoid prohibiting any activity or approach that is required by

those standards. In particular, compatibility with ISO 26262:2018 and ISO/PAS 21448:2019 has been considered. Annex A discusses a mapping of some clauses of this standard onto ISO 26262:2018 and ISO/PAS 21448:2019. Other safety standards such as IEC 61508 are relevant and expected to be generally compatible, but detailed analysis of IEC 61508 and other functional safety standards is out of scope for this version of UL 4600.

1.3.4 Two areas out of scope for this standard are setting acceptable risk levels and setting forth requirements for ethical product release decisions and any ethical aspects of product behavior. For both topics the developer records what decisions have been made, but this standard does not establish acceptance criteria beyond that they have been recorded. Other standards such as the IEEE P7000 series provide guidance on those topics.

2 Scope

2.1 Scope summary

2.1.1 This standard covers the safety principles, risk mitigation, tools, techniques, and lifecycle processes for building and evaluating a safety argument for vehicles that can operate in an autonomous mode.

2.1.2 Operation is assumed to occur without human supervision and without expectation of human intervention in performing and supervising the dynamic driving task and other normal system operations based upon the current item state and ability to sense and otherwise interpret the operating environment. Human contributions to safety in other than normal operation are considered (e.g., maintenance), as are interactions with humans who are not operating the item (e.g., pedestrians).

2.1.3 This standard generally uses the term “item” rather than “system” or “product” when referring to the scope of the safety case as well as the operation of the item. This approach is in recognition of the possibility that the safety of the item might rely upon infrastructure, services, support processes, and other factors that might not normally be considered part of a system such as a vehicle per se, but which materially affect its safety and therefore are all considered within the scope of the item being assessed for conformance.

2.1.4 This standard assumes that the item autonomously operates starting at some well-defined initial state to some other well-defined end state without human intervention. Human input might influence the selection of desirable states (e.g., via an occupant requesting a destination). However, the extent to which human operators mitigate or introduce risk by performing or supervising a dynamic control task (e.g., by driving or taking responsibility for monitoring system operation) is outside the scope of the standard. Similarly, the extent to which human operator performance or non-performance is involved in risks related to transferring human driver control to or from the item is also outside the scope of the standard. However, ensuring that the item itself properly performs any change of control functions if and when it is supposed to is generally within the scope of the standard since it can adversely affect operation in fully autonomous mode as well. Thus, while portions of this standard might be helpful for addressing less than fully autonomous vehicles, issues involving human driver responsibilities, vigilance, and ability to properly accept responsibility for vehicle control are out of scope for this standard.

2.1.5 While information security is an essential topic, the details of that area are out of scope for this standard beyond a general requirement for a Security Plan and prompt elements that are possibly unique to autonomous vehicle operation in comparison to other vehicular security requirements. Reasonably foreseeable misuse and abuse as well as physical attacks (e.g., physical sensor damage) are in scope.

2.1.6 The requirements of this standard are considered to be at a necessary, but possibly not sufficient, level of completeness and rigor to create an acceptably well-formed and acceptably complete item safety

case. In particular, prompt element lists are considered non-exhaustive, with an expectation that design teams will include additional items as relevant to the item and its operational design domain.

2.2 Elements in scope

2.2.1 Specific aspects of item operation and safety related issues which are explicitly intended to be in scope for this standard include:

a) Operation of autonomous items in potentially unstructured environments

EXAMPLE: A vehicle is the first vehicle directed into an open farm field containing a mixture of viable and non-viable areas for traversal and parking as part of an ad hoc overflow event parking process. There are no lane markings and no positioning beacons. Moreover, there are cows and hay bales randomly placed in the field. There are no humans assisting with organizing vehicle parking positions. This situation is in scope for the standard.

EXAMPLE: A crowd has spilled into the street at a fire scene. Emergency response equipment, response personnel, victims, and casual observers are moving without regard to normal road use patterns. Fire hoses, falling pieces of burning debris, small explosions, traffic signal power outages, damaged pavement, and other disruptions to normal infrastructure expectations exist. Multiple injured people at building exits are calling for pickup by autonomous vehicle ride hail services to be transported to urgent care medical facilities. This situation is in scope for the standard.

NOTE: A particular item's safety case might require a structured environment for safe operation as specified by an ODD description. However, structure is not assumed to be present by default. Therefore, operation in unstructured environments must be specifically disclaimed by the safety case if applicable.

b) Operation with potentially inaccurate, incorrect, incomplete, or misleading data provided by sensors

c) The effects of potentially inaccurate, incorrect, incomplete, or biased data, including test data, field report data, other validation data and machine learning training data.

d) The effects of potentially imprecise, inaccurate, or incomplete simulation models.

e) Potential defects and failures of hardware and/or software in the item, data collection functions, data processing functions, communications, engineering support systems, tools, and infrastructure support.

f) Human contributions to potential risk, including occupants, pedestrians, other road users, non-road users, cargo handlers, maintainers and inspectors. This includes acts of omission and commission; accidental and malicious physical acts; and human roles in creating as well as mitigating risk.

g) Lifecycle considerations, including design data collection, engineering data management, tool qualification, design, implementation, testing, other validation, field data collection, operations, maintenance, updates, upgrades, and retirement. Lifecycle considerations also encompass potential changes to the environment which may affect ODDs, changes in object types, changes in behaviors, etc..

h) Inclusion of risk mitigation and other aspects of contributions to the safety case made by conformance with other standards, and in particular both ISO 26262 and ISO/PAS 21448 standards for products within scope for those standards.

- i) Ability to use a heterogeneous approach to arguments, including use of diverse standards to support safety (e.g., use of different but acceptable functional safety standards for different item subsystems).

2.2.2 None of the described in-scope topics is intended to require that the item successfully delivers full service in all situations described. Rather, the requirement is to consider all prompt elements and argue that risk is acceptable despite these factors. In many cases that will involve crafting an ODD that excludes problematic prompt elements. However, excluding a prompt element from the ODD (or similar approach) creates an obligation to argue that the exclusion does not itself result in unacceptable risk.

EXAMPLE: Unpaved roads without lane markings are excluded from the ODD. The safety case generally argues that geo-fencing and map creation will exclude all unpaved roads. It is further argued that this exclusion encompasses quickly identifying roads undergoing repaving projects that are temporarily unmarked but still carrying traffic.

EXAMPLE: Snow is excluded from the ODD. Snow is still part of the safety case to cover un-forecast snow that occurs during an operational mission. The safety case generally argues that it can successfully terminate a mission via in-lane stop despite snow. It further argues (with evidence) that snow will happen so infrequently in the deployment location that the elevated product risk presented by occasional in-lane stops is acceptable.

2.3 Scope limitations

2.3.1 A significant scope limitation of this standard is that it does not cover the detailed topics relevant to ensuring that humans are able to provide effective safety supervision for an autonomous item. Rather, coverage is limited to fully autonomous operation with no human supervision as well as aspects of the item during human supervised operation that do not relate to arguing human supervision effectiveness. Similarly, aspects of the ability of a human operator to safely control the item are out of scope. More specifically, the following are explicitly intended to be out of scope for this standard:

- a) Aspects of items as well as item-level safety of operational modes for which the locus of control is outside the item itself

EXAMPLES: End-to-end system-level safety (including the human supervisor or driver) of teleoperated modes of operation is excluded to the degree it relies upon a human teleoperator to control, supervise, or otherwise ensure system safety.

NOTE: A product-level “item” is intended to include offboard functions that participate in control of a vehicle, such as a cloud-based route planning system.

NOTE: The proper response to teleoperated commands and proper transmission of teleoperation data out of the vehicle is in scope. However, whether teleoperation is actually safe at a system level is out of scope due to human involvement in the driving task.

- b) Human factors related to safety during or after a handoff or mode switch that makes a human responsible for the dynamic control task safety.

EXAMPLE: The details of ensuring that a human supervisor is available and able to safely take over item operation upon request and continue to operate the vehicle safely when some or all autonomy functions are disabled are out of scope.

- c) Risk mitigation or other safety argument credit taken for the contribution of humans to the dynamic driving task (e.g., human driver, human safety supervisor, human teleoperator, issues of human alertness, issues of human situational awareness).

EXAMPLE: The details of how safe and effective human/machine interfaces should be provided to teleoperators are out of scope.

d) Road testing of prototype vehicles to the degree that safety argument takes credit for a human performing and/or supervising the dynamic driving task.

e) Specifics regarding how to ensure that humans acceptably meet expectations for non-driving roles in item safety. To be clear, identifying the human contribution to the safety case (e.g., via performing inspections, or a human's ability to correctly perceive and interpret signals provided by the item) is in scope. However, specifying the details of how to actually ensure that the contribution is being done in an acceptable manner in terms of human behaviors, psychology, limitations, and so on is out of scope. While competency frameworks, staff skill lists, and experience requirements are potentially helpful topics to cover a safety case, specifics regarding these topics are out of scope for this standard.

f) Specifics regarding how to evaluate the suitability and effectiveness of human interface devices. To be clear, the need to identify such devices and the need to ensure suitability and effectiveness is in scope, but specifying requirements for how to meet that need is out of scope.

EXAMPLE: A car in autonomous operation is not supposed to transfer vehicle control to an occupant under any circumstance during a mission. It does in fact attempt to transfer vehicle control to an occupant, providing three seconds of warning.

In Scope for UL 4600: Vehicle attempting a handoff to a driver when it was not supposed to.

In Scope for UL 4600: Incorrectly attempting to transfer control in violation of fully autonomous operational mode.

Out of Scope for UL 4600: Whether three seconds is enough warning for effective handoff and whether the occupant is a qualified driver.

EXAMPLE: An autonomous car is designed to transfer control to a human driver under some circumstances with a 10 second warning when the driver is qualified, competent, and aware of this handoff mission parameter.

In Scope for UL 4600: Transferring control without the full 10 seconds of warning; A defectively designed brake control mechanism that under some circumstances prevents the human driver from actually regaining control at the designated time (e.g., human driver's brake pedal disabled despite having attempted to perform a handoff); Whether brake pedal actuation by a human is intended to initiate a handoff under specified conditions.

Out of Scope for UL 4600: Whether ten seconds is enough warning for effective handoff in any particular handoff situation; Whether depression of a brake pedal by a human driver is a safe (from human factors point of view) handoff initiation mechanism; Checking whether occupant is a qualified driver; Checking whether the driver has sufficient cognitive ability for safe operation; Checking whether human driver is in correct seating position to assume operational control; Safety of vehicle once control has been transferred to the human driver; Effectiveness of Advanced Driver Assistance (ADAS) functions in mitigating risk while under human driver control.

2.3.2 There are a number of additional topics out of scope. Reference to these topics should be made where relevant to the safety case, but specifics such as prompt elements to provide technical depth are not included in this standard:

a) The specific intended function (e.g. surface cleaning, fragile cargo delivery)

NOTE: This topic might be covered by an end product standard.

b) End product requirements

NOTE: It is intended that end product standards can reference or require this standard.

c) Legal and policy issues

EXAMPLES: Determining liability, what records retention policies are appropriate, what level or product risk is actually acceptable to society.

d) Ethical issues

EXAMPLES: Resolving questions of acceptable risk, evaluating comparative severity of different loss event scenarios

e) Electric Vehicle safety

EXAMPLES: Safe battery design, safe battery management algorithms, battery thermal management

f) General vehicle safety

EXAMPLES: Crash mitigation, occupant restraints, refueling/recharging safety

g) Non-safety related quality aspects of performing the intended function

EXAMPLES: Ride quality, fuel economy

h) Effectiveness of crash and injury mitigation mechanisms

EXAMPLES: Seat belts, air bags, child seats

2.3.3 Also out of scope for the standard is any implied redefinition of existing standards and accepted practices to the degree that they are acceptable to support the safety case being made. Reference to these topics should be made where relevant to the safety case, but specifics are not included in this standard. These topics include:

a) Government regulations

EXAMPLES: FMVSS, Federal Communications Commission radio frequency interference emission certification

b) Mechanical aspects of the item

EXAMPLES: Sharp edges, pinch points, window lift motor closing pressure limits

c) Legacy operational procedures

EXAMPLES: Car door operation, child locks on car doors, cargo loading and loading

d) Specifics of vehicle support for other than normal operations (except to the degree that the autonomy is expected to provide such support and that such provision is safety related)

EXAMPLES: Autonomy support for non-safety related routine maintenance procedures in situations which do not present a hazard

e) Sufficiency and performance of controls and item response when a human is operating or supervising the dynamic driving task

EXAMPLE: As required by FMVSS

f) Licensing, training, qualification and other aspects of ensuring human competence to operate or participate in the vehicle lifecycle

NOTE: Credit can be taken for these in the safety case only to the degree that any standard or procedure conferring or documenting qualification provides objective evidence of abilities

3 Referenced Publications

3.1 Normative references

3.1.1 The following referenced documents are indispensable for the application of this document:

N/A

3.2 Informative references

3.2.1 The following references included in this standard are for information only. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

Def Stan 00-56, *Safety Management Requirements for Defence Systems*

EN 50129:2018, *Railway applications – Communication, signaling and processing systems – Safety related electronic systems for signaling*

FAA AC 25.1309-1A, *System Design and Analysis*

FAA FO 8040.4B, *Safety Risk Management Policy*

FAA FO 8110.49A, *Software Approval Guidelines*

IEC 61508-4, *Functional safety of electrical/electronic/programmable electronic safety-related systems – Part 4: Definitions and abbreviations*

IEEE 1012-2012, *System, Software, and Hardware Verification and Validation*

IEEE 24765, *Systems and Software Engineering – Vocabulary*

IEEE P7000, *Standard for Model Process for Addressing Ethical Concerns During System Design*

ISO 26262:2018, *Road vehicles – Functional safety*

ISO/PAS 21448:2019, *Road vehicles – Safety of the intended functionality*

MIL-STD-882E, *Standard Practice System Safety*

Reinventing Safety: A Joint Approach to Automated Driving Systems (Mercedes Benz & Bosch):
<https://www.daimler.com/documents/innovation/other/vssa-mercedes-benz-and-bosch.pdf>

Safety First for Automated Driving; Published by eleven automotive and mobility industry leaders – a first-of-its-kind framework for safe automated passenger vehicles: <https://newsroom.intel.com/news/intel-auto-industry-leaders-publish-new-automated-driving-safety-framework/#gs.dr3a1>

SCSC-153A, *Safety Assurance Objectives for Autonomous Systems* – Safety Critical Systems Club – UK;

Note: this document was written by the Safety of Autonomous Systems Working Group (SASWG), which is convened under the auspices of the SCSC, UK. The goal of the SASWG is to produce clear guidance on how autonomous systems and autonomy technologies should be managed in a safety related context, throughout the lifecycle, in a way that is tightly focused on challenges unique to autonomy. The document intends to address safety due to Artificial Intelligence (AI) from Machine Learning only and is scheduled to be formally released at SSS'20 11-13 February 2020. Link to document: <https://scsc.uk/SCSC-153A>.

4 Terms, Definitions, and Document Usage

4.1 How to interpret normative elements (Normative)

4.1.1 Commonly used constructions of this standard affect the safety case as follows. All elements are normative except “EXAMPLE,” and “REFERENCE” statements as well as any other content that is explicitly stated to be informative. (See [Table 4.1](#) below for a summary of key safety case deviation explanations.)

a) **Numbered clauses** (starting at [5.1.1](#)) are generally stated as “shall” conformance obligations. These are intended to be general statements, with supporting normative prompt elements providing further detail. Each clause is specifically addressed in the safety case with the exception of conformance assessment process clauses in Section [17](#) that deal with activities performed upon the safety case itself. An important part of navigability of the safety case is a capability to identify the portion(s) of the safety case that support fulfillment of each clause. The scope of all clauses is the safety related portion of the item unless otherwise stated.

b) **MANDATORY prompt elements:** Addressed by the safety case. Safety case deviations not permitted. Any safety case deviation results in a non-conformance.

EXAMPLE: “Identify hazards” is mandatory – it must be done.

EXAMPLE: A team attempts to argue that MANDATORY prompt element X does not apply to their item. This is an invalid attempt at a safety case deviation.

NOTE: In some cases, a MANDATORY prompt element refers to consideration of a different clause in a hierarchical manner. That should be interpreted as a mandatory inclusion of the associated higher-level claim in a safety argument, but not mandatory inclusion of all the non-mandatory prompt elements of the clause being referred to. In particular, such hierarchical references are not intended to override the safety case deviation rules.

EXAMPLE: MANDATORY prompt element X states that section Y is addressed by the safety case. Section Y has a HIGHLY RECOMMENDED prompt element Z. The net requirement is that satisfaction of all clauses in Section Y must be addressed by the safety case, but a safety case deviation of prompt element Z is still permitted in accordance with its HIGHLY RECOMMENDED categorization.

c) **REQUIRED prompt elements:** Addressed by the safety case. Safety case deviation is permitted only if documented by argument that the prompt element is intrinsically incompatible with the item and/or its safety case. Support for each safety case deviation is explicitly noted in the safety case. End product standards can enumerate REQUIRED elements that can be omitted from the safety case (i.e., blanket default safety case deviations specified by an end product standard). The safety case is non-conformant if safety case deviations are not acceptably documented for REQUIRED elements. Safety case deviation for a reason other than intrinsic inapplicability results in a non-conformance. Field engineering feedback and change impact analysis are used to detect the possibility of a claim of intrinsic incompatibility becoming invalid. Examples of acceptable safety case deviations include:

EXAMPLE: Safety case deviation for requirements on machine learning if item does not use machine learning-based techniques in any manner, including design, operation, and field operational data analysis.

EXAMPLE: Safety case deviation from recognizing road signs if item does not rely upon road signs in any safety related way. Arguments supporting this might be that the ODD specifically excludes road signs, or that the item exclusively uses a means other than road signs to gather equivalent information.

EXAMPLE: Safety case deviation from requirements specific to subdivision of an ODD into multiple ODD subsets if the item does not define ODD subsets (i. e., if an item uses a single monolithic ODD, then any requirement related to ODD subsets is not applicable).

EXAMPLE: Data does not yet exist because a potential event or condition has never occurred in the life of the item. An example would be records of correction of defects discovered during item operations if no items have yet been deployed. However, safety case deviation is not permitted for the mechanisms and procedures to collect and process such potential events – just the portions of the safety case contents that cannot exist until field events occur. It is a good practice to include clearly designated placeholder examples to populate empty data sets so as to exercise tools, procedures, and traceability.

d) **HIGHLY RECOMMENDED prompt elements:** These are best practices that should be followed, but may be omitted, especially for low risk items. Omissions are explicitly noted in the safety case with reasonable supporting argument to provide a hook for tracing root cause analysis back to those omissions. The safety case is considered acceptable so long as these omissions are noted with an acceptable rationale. In cases in which a generic prompt element of “others” is included, an omission rationale of “no others” is acceptable. A primary purpose of field engineering feedback is to ensure that a safety case deviation of any prompt element that contributes substantively to safety related issues is identified and the deviation revoked.

EXAMPLE: The use of a specific analysis technique is HIGHLY RECOMMENDED. The safety case notes that the technique was not used with a rationale of “other analysis techniques being used provide comparable information.” The safety case includes an argument that there is no history of incidents with an unknown root cause that could plausibly have been prevented via addition of this technique to the design approach. This argument is backed up by root cause analysis logs showing all root causes have been resolved, leaving no unresolved candidates that might in fact trace to the prompt element deviation in the absence of other adverse information.

EXAMPLE: The rationale for a HIGHLY RECOMMENDED prompt element deviation is “Safety case review meeting of 9/23/2019 determined this was inapplicable.” While light on technical content, this is specific documentation that a deliberate process was said to be used to decide not to address a prompt element. However, root cause analysis must still be performed of field issues, and might potentially invalidate this rationale depending upon field experience.

e) **RECOMMENDED prompt elements:** These are optional prompt elements documenting good practices and/or suggestions for helpful techniques. If adopted, they can be included in the safety case. However, addressing these prompt elements is entirely optional. The safety case is considered well-formed whether they are included or not.

f) **PITFALL prompt elements:** These are anti-patterns, typically of the general form: “If X is true, then item is prone to increased risk of Y.” The intended interpretation is that if X is true (e.g., use of some design pattern X or engineering technique X appears in the safety case), then the safety case is considered invalid unless epistemic defeater Y has been explicitly argued to be false. In other words, Y presents item risk that is presumed to have been activated by X unless the safety case presents reasonable argument and evidence of mitigation of risk Y. More formally, the term “Pitfall” tags a conditional epistemic defeater prompt regarding defeasibility of a claim that the parent requirement has been met. Responding to a Pitfall prompt element might be a claim that

in two ways. (1) Arguing that the precondition X is not true. This amounts to an acceptable safety case deviation for the prompt element. (2) Arguing that risk due to post condition Y is mitigated when or if precondition X is or might be true. Safety case deviation rules for Pitfalls apply according to the categorization (MANDATORY, REQUIRED, HIGHLY RECOMMENDED, or RECOMMENDED). A safety case deviation has occurred when there is no argument that the Pitfall has been avoided.

NOTE: For assessment purposes each Pitfall only applies to the scope of the specific clause in which it is listed. However, it is a good practice (but optional) to consider while creating the safety argument that Pitfalls might have a larger impact.

g) **CONFORMANCE statements.** Conformance with each clause is evaluated via both self-audit and independent assessment according to Section 17. Each clause has a conformance statement that provides guidance identifying portions of the safety case and other information sources that are especially relevant to assessing conformance to that clause. Assessors are permitted to consider objective evidence beyond the conformance statement when the assessor determines that the situation warrants, but are limited in conformance determination by the written scope of the clause. (Self-auditors can and should consider whether prompt elements beyond those included in this standard need to be added for a particular item's safety case.) Conformance checks are performed by someone other than the developer of the item being checked for conformance (See Section 17). (There are limited exceptions for safety case artifacts being self-audited; see Section 17.2.2.)

h) **NOTE statements.** These are notes on how to interpret the normative elements of a section and in some cases provide rationale material.

i) **EXAMPLE lists (non-normative).** In some cases examples are provided. Examples are non-normative and, in many cases, will not apply to all items. Their primary purpose is to define by example to reduce potential ambiguity. A secondary purpose is to serve as an informative but incomplete checklist of commonly occurring topics that should be considered if applicable. It is expected and required that construction and assessment of the safety case go beyond the bounds of any examples. While exclusion of examples in a safety case does not by itself result in a finding of non-conformance, independent assessors are encouraged to suggest additional examples, as well as elevation of specific examples to normative prompt element status when warranted based on assessment experience.

j) **REFERENCE (non-normative).** References provide citations to materials such as other standards. They are non-normative by default.

4.1.2 Summary of safety case deviation approaches for elements of different types is shown in [Table 4.1](#):

Table 4.1
Safety Case Deviation Approach

MANDATORY	<u>No safety case deviations permitted.</u>
REQUIRED	<u>Safety case deviations only for requirements that are intrinsically inapplicable</u> due to the fundamental nature of the item and/or the current deployment state of the item. All safety case deviations recorded in safety case with justification. Impact analysis and lifecycle tracking monitor the possibility of a change of applicability status.
HIGHLY RECOMMENDED	<u>Safety case deviations permitted with an acceptable rationale.</u> Impact analysis and lifecycle tracking monitor the possibility of a change of applicability status. All safety case deviations recorded in safety case with justification.
RECOMMENDED	<u>Optional items.</u> Need not be mentioned by safety case. No argument support required for safety case deviation.

4.1.3 Lists that support a particular clause, such as a list of MANDATORY or REQUIRED items, are interpreted in the following manner:

- a) For the sake of brevity, the context of each item in a list is the overarching clause, even if not explicitly stated. This means lists generally provide a set of prompts for brevity and usability rather than fully stated “shall” statements.
- b) There is no implied ranking, preference, or priority implied by list order unless explicitly stated.
- c) **All prompt element lists are presumed to be potentially incomplete unless specifically stated otherwise. A well-formed and complete safety case extends lists as needed to achieve an acceptable safety outcome.** Unless otherwise stated, prompt element lists are to be considered as a minimum set of points to be considered in the safety case, and not an exhaustive set of all possible points.
- d) If different prompt element lists appear to overlap, it is acceptable to trace a single argument or evidence branch to multiple prompt elements if desired. A common situation is that a MANDATORY prompt element will give a generalized prompt element, whereas one or more REQUIRED prompt elements will give specific examples of that generalized element to ensure that they are considered if applicable.
- e) If a single prompt element list contains apparently overlapping items, it is sufficient to trace to whichever prompt element seems most applicable to the developers. Such overlapped lists are often used to address issues of potentially different interpretations and terminology across domains and sub-domains.
- f) Each prompt element in the list is accounted for at the stated level of safety case deviation rigor on a per-element basis. Therefore, if safety case deviation is desired for a ten-element list for a single REQUIRED section, the safety case documents safety case deviation in a way that is individually traceable to each of the ten prompt elements in the list. (A single justification statement for safety case deviation might trace to some or all of the ten elements, but that traceability must be documented.)
- g) “At least one of” phrasing contains a sub-list of alternatives. Only one alternative within the list need be addressed in the safety case, although others might additionally be addressed if desired. This is true even if the “at least one of” occurs in a MANDATORY or REQUIRED list. Note that in some cases two or more prompt elements might be required in practice due to Pitfall statements. Alternatives not listed can be considered to meet the “at least one of” criterion with suitable argument support as to acceptability. (This is according to the principle that prompt element lists are not considered to be exhaustive, and therefore safety cases can add prompt elements as appropriate to the local version of those prompt element lists.)
- h) “Tailoring” of list items must adhere to safety case deviation rules summarized in [Table 4.1](#) above.

4.1.4 For simplicity and uniformity, references to sections, clauses, and prompt elements are made using a decimalized notation, notwithstanding the typographical conventions used in numbering the sections themselves.

EXAMPLE: 9.2.3.3.b.1.i is a correctly formatted reference to a hypothetical prompt element sub-sub-list in the HIGHLY RECOMMENDED category of Section 9.2.3.

4.1.5 See also [2.1.6](#) regarding scope of prompt lists.

4.2 Terms and definitions (Normative)

4.2.1 Acceptable

Sufficient to achieve the overall item risk as determined in the safety case.

Example: “Acceptable test coverage” means that the amount of test coverage is sufficient to support the safety case’s claim of overall item risk after taking claimed risk mitigation credit.

NOTE: This is an objective term related to the validity and completeness of the safety case, and not a subjective term related to any particular assessor’s personal point of view.

See: Section [6.4.3](#)

4.2.2 Activation (of faults or hazards)

An input or situation that causes the system to potentially fail due to a fault or hazard.

EXAMPLE: A bit in memory corrupted by a single event upset is a fault. That fault is activated when the memory location is read and results in a computational error that causes a failure in the form of an incorrect or unsafe item behavior.

NOTE: Fault mitigation such as error detection coding, among other mitigations, can prevent an activated fault from causing a failure.

4.2.3 AI Techniques

A general description of computational algorithms and other techniques that include inductive learning, intentionally non-deterministic behavior, rule-based systems, computer vision, heuristic searches, and other techniques. These are often referred to as “Artificial Intelligence” techniques.

NOTE: This term is meant to be a broadly interpreted descriptive term to encompass software that is not generally amenable to pre-AI software safety approaches. Whether or not actual “intelligence” is actually involved is a matter beyond the scope of this standard.

4.2.4 Argue

Construct a safety case (including claims, arguments, and evidence) that a particular requirement has been met. The resulting argument shows that evidence supports the claims.

EXAMPLE: “Developer shall argue that all hazards have been mitigated” directs the developers to include in the safety case claims, arguments, and evidence that each hazard has in fact been mitigated.

4.2.5 Assessor

One or more people who perform assessment.

4.2.6 Assessment

Review and evaluation of the safety case.

NOTE: Two types of assessment are encompassed by this standard: self-audits and independent assessment.

4.2.7 Automated

Autonomous.

4.2.8 Autonomous

Operates without human oversight or intervention.

NOTE: This standard treats the terms “automated” and “autonomous” as generally interchangeable.

4.2.9 **Autonomy**

A capability or function that provides autonomous operation.

4.2.10 **Brittleness (of a system)**

Tendency to fail when encountering conditions that are only slightly different from conditions in which it works properly.

4.2.11 **Chaotic (of a system)**

Sensitive to initial conditions to the point of appearing to exhibit random behavior.

EXAMPLE: A robot that is pointed exactly at the center of mass of an obstacle using a completely deterministic algorithmic approach might decide on an apparently random basis to go around the obstacle to the left or right depending on minute variations in conditions beyond the capability of a tester to accurately control, making it impossible in practice to determine which action it will take.

4.2.12 **Claim**

A falsifiable statement that contributes to establishing that safety is acceptable.

4.2.13 **Conformance**

An independent assessment result indicates that the item’s safety case meets the requirements of UL 4600. (See Section [17.1](#))

4.2.14 **Criticality Level**

A level categorizing the risk associated with an unmitigated hazard.

NOTE: Criticality level is intended to be a generic term that encompasses integrity level approaches and assurance level approaches such as SIL, DAL, and ASIL.

4.2.15 **Demonstration**

Operation of the item, functions, or elements to show that a specific property, functionality, or other aspect of the item matches the safety case.

4.2.16 **Doer/Checker**

A heterogeneous architectural design pattern in which a Doer Fault Containment Region (FCR) performs a function while a second Checker FCR performs acceptance tests on the Doer, performing a mutual shutdown if the acceptance test fails.

NOTE: This general concept is also known in other contexts as a safety bag or monitor/actuator pair. It is a specific strategy for using an acceptance test.

4.2.17 **Evidence**

Data or other information used to support arguments.

4.2.18 **Fault Containment Region (FCR)**

A collection of elements “that operates correctly regardless of any arbitrary logical or electrical fault outside the region.”

Reference: (Lala, J., and Harper, R. Architectural principles for safety-critical real-time applications, Proceedings of the IEEE, 82(1), Jan. 1994, pp. 25-40.)

NOTE: Two FCRs are required to ensure fault detection and/or fault mitigation in the presence of an arbitrary FCR failure.

4.2.19 **Fault Model**

A specification of all the faults (and their types) that are considered when performing fault analysis.

NOTE: A fault model can include multiple concurrent faults of different types.

4.2.20 **Field Engineering Feedback**

Acquisition of data from system operation for the purpose of supporting the safety case and identifying potential safety case issues.

4.2.21 **Identify**

Create an enumerated list responsive to a specified category, property, or other aspect of a clause with sufficient specificity to enable assessment.

(Alternate word form: Identification)

4.2.22 **Incident**

Occurrence of a safety-related failure which might have resulted in a loss event

NOTE: An incident does not necessarily result in a loss event. It is sufficient that in other circumstances an incident might possibly have resulted in a loss event.

EXAMPLE: A car fails to stop at a stop sign. There is no cross traffic, so no collision results. If cross traffic had been present, a collision could have occurred. This is an incident even though no loss event occurred.

NOTE: All loss events are also incidents. Therefore, the phrase “incidents” is equivalent to “incidents and loss events.”

4.2.23 **Independent (failure)**

Two or more Fault Containment Regions (FCRs) that have no correlated, common cause, and/or common mode fault conditions.

NOTE: Assuming no accumulation of faults over time, the probability of simultaneous fault activation can be expressed by a simple product of unconditional probabilities.

4.2.24 **Independent (review and assessment)**

assessor or reviewer who has no direct incentive and no substantive indirect incentive influencing the result of a review or assessment process (see also [17.3.2](#)).

4.2.25 **Item**

A product, element, system of systems, or other product-related scope for which conformance to UL 4600 is assessed.

NOTE: The “item” may need to include infrastructure, offboard computing, offboard data storage, development processes, lifecycle support processes, supply chain quality assurance measures, and other aspects of ensuring safety beyond the boundaries of a deployed product itself. The “item” can include an entire product, or only portions of a product, but in any event will include all aspects required for conformance of the portion of the product contained within the item.

4.2.26 **Life Critical**

An aspect of a system for which a loss event could potentially result in the loss of a human life

NOTE: This is strictly a severity concept. A hazard can be life critical even if developers argue that the associated risk is low due to an improbability of occurrence in operation.

4.2.27 **Loss**

Human death, human injury, animal death, animal injury, property damage, environmental damage, or other substantive adverse outcome.

NOTE: Which losses are considered adverse is system specific. However, human death and significant human injury are always to be considered to be adverse losses.

NOTE: Safety cases might elect to consider a financial cost caused by an item failure to also be a “substantive adverse outcome.”

NOTE: “Loss event” generally corresponds to the term “accident” in FAA Order 8040.4B and DefStan 00-56. However, the term “Loss” is used to avoid any preconceptions regarding liability and foreseeability.

4.2.28 **Mitigate**

Reduce to an acceptable level of risk.

NOTE: The level of acceptable risk depends upon the criticality of the specific risk being mitigated (e.g., life-critical vs. not life critical) and its contribution to item-level risk as established in the safety case. A hazard that presents acceptably low risk without overt mitigation action is still considered “mitigated” so long as it is argued that the risk presented is acceptably low. The word “mitigated” is equivalent to the concept of “acceptably mitigated.” (See definition of “acceptable.”) A risk is not considered mitigated until the mitigation approach, if any, is implemented and associated safety case arguments regarding mitigation have been updated (“tracked to closure”). An “accepted risk” is one for which partial (or no) mitigation has been performed (See Section [5.5.1](#))

4.2.29 **Nondeterminism**

Behavior that is not repeatable.

NOTE: Nondeterminism might be caused by real time scheduling perturbations, use of pseudo-random algorithms, or other factors.

See also: Chaotic

4.2.30 **Operational Design Domain (ODD)**

The set of environments and situations the item is intended to operate within. This includes not only direct environmental conditions and geographic restrictions, but also a characterization of the set of objects, events, and other conditions that will occur within that environment.

NOTE: A system has a single ODD by definition. Assessment is made with regard to the entire ODD.

See also: ODD Subset

4.2.31 **ODD Subset**

A managed portion of an item’s ODD.

EXAMPLE: An all-weather ODD is broken up into subsets for fair weather, rain, and snow/ice.

NOTE: An ODD subset might be defined to partition the operational space to ease design tasks, support phased deployment by adding additional subsets over time, or otherwise manage the complexity of a

potentially large and varied ODD. The safety case might argue each ODD subset independently for some aspects of the safety case.

4.2.32 Proof Test

A test that goes beyond Built In Self-Test (BIST) capabilities to detect latent faults.

NOTE: Proof tests historically referred to mechanical tests such as pressure testing vessels or moving emergency valves that cannot be done during normal system operation. They often involve exercising failsafes, exercising sensors that are used to detect failures, or otherwise performing off-line tests.

4.2.33 Provenance

The background of components, materials, supplies, software, and other aspects of the item, especially as conferring distinction or quality.

NOTE: In the context of UL 4600, this refers to establishing that COTS products and their components actually provide the required capabilities, and specifically excluding inferior, counterfeit, and other “unapproved” parts. This is different than variations in parts that provide acceptable capabilities and other versioning activities.

4.2.34 Quality Fade

A supply chain fault in which a component supplier degrades component quality over time via progressive use of substitute materials, substitute components, design changes, and/or elimination of protective components while apparently maintaining functionality and meeting tested parameter values.

4.2.35 Risk

A combination of the probability of occurrence of a loss event and the severity of that loss event.

NOTE: Risk is typically some weighted combination of probability and severity, potentially with a zero or non-linear weighting given to probability. This definition is not meant to exclude alternate but comparable formulations of risk.

4.2.36 Robust

Able to continue operation despite situations that are out of specification.

EXAMPLES: System continues normal or degraded operation despite receiving out of specification inputs, encountering ODD violations, suffering component faults, and encountering data not represented in machine learning training sets.

NOTE: Robustness is often a matter of degree rather than an absolute property.

4.2.37 Safe

Having an acceptable post-mitigation item level risk as defined by the safety case.

See Section [6.4.3](#).

4.2.38 Safety Case

A structured argument, supported by a body of evidence, that provides a compelling, comprehensible and valid case that a system is safe for a given application in a given environment.

Source: Defence Standard 00-56 Issue 7 (Part 1): Safety Management Requirements for Defence Systems. UK Ministry of Defence. p. 26.

4.2.39 Safety Case Deviation

A recorded exclusion of a prompt element from consideration in accordance with the requirements of Section [4.1](#).

4.2.40 **Safety Related**

A function, element, property, or other aspect of the system that affects safety directly or indirectly.

4.2.41 **Element out of Context (EooC)**

A stand-alone element that is the subject of a safety case fragment and assessment.

See Section [5.7.3](#).

4.2.42 **Safety Performance Indicator (SPI)**

A metric used to quantify safety performance.

NOTE: This term is analogous to the term Key Performance Indicator (KPI), but is specific to safety related aspects of the item.

GENERAL TERMINOLOGY REFERENCES (INFORMATIVE):

- Avizienis, A., Laprie, J-C., Randell, B., Landwehr, C., "Basic Concepts and Taxonomy of Dependable and Secure Computing," IEEE Trans. Dependable and Secure Computing, 1(1), Jan.-Mar. 2004, pp. 1-23.
- ISO 26262-1:2018
- ISO/PAS 21448:2019

4.3 Abbreviations and Acronyms (Informative)

- a) AI: Artificial Intelligence
- b) ASIL: Automotive Software Integrity Level
- c) BIST: Built-In Self-Test
- d) COTS: Commercial Off The Shelf
- e) DAL: Design Assurance Level
- f) DTC: Diagnostic Trouble Code
- g) ECU: Electronic Control Unit
- h) EooC: Element out of Context
- i) FCR: Fault Containment Region
- j) FMVSS: Federal Motor Vehicle Safety Standards (USA)
- k) GNSS: Global Navigation Satellite System
- l) GPU: Graphics Processing Unit
- m) HARA: Hazard Analysis and Risk Assessment
- n) ISO: International Standards Organization

- o) MEL: Minimum Equipment List
- p) NDI: Non-Development Item
- q) ODD: Operational Design Domain
- r) PSSA: Preliminary System Safety Assessment
- s) SEBOK: System Engineering Body of Knowledge
- t) SIL: Safety Integrity Level
- u) SOUP: Software of Unknown Provenance / Systems of Unknown Provenance
- v) SPI: Safety Performance Indicator
- w) SQA: Software Quality Assurance
- x) SWEBOK: SoftWare Engineering Body of Knowledge
- y) V&V: Verification and Validation
- z) V2I: Vehicle-to-infrastructure
- aa) V2V: Vehicle-to-vehicle
- bb) V2X: Vehicle-to-other

5 Safety Case and Arguments

5.1 General

5.1.1 The safety case shall be a structured explanation in the form of claims, supported by argument and evidence, that justifies that the item is acceptably safe for a defined operational design domain, and covers the item's lifecycle.

5.1.1.1 MANDATORY:

- a) Conformance is demonstrated based upon a documented safety case that:
 - 1) Uses an acceptable safety case format (see Section [5.2](#))
 - 2) Presents an acceptably complete argument arguing that the provided evidence supports the defined claims (see Section [5.3](#))
 - 3) Presents acceptable evidence (see Section [5.4](#))
 - 4) Addresses accepted risks (see Section [5.5](#))
 - 5) Addresses safety culture (see Section [5.6](#))
- b) Configuration management of the safety case

5.1.1.2 REQUIRED:

- a) Addresses Elements out of Context (EooC) (see Section [5.7.3](#))
- b) Additional evidence not included in the safety case is provided upon request to an independent assessor or self-auditor

EXAMPLES: Test results are summarized in the safety case. Details and any additional descriptive material of evidence are made available upon request.

c) Any aspect of a purported safety case that is not documented or not provided is disregarded in performing the assessment.

d) Unless an Element out of Context (EooC) assessment is being performed, the safety case encompasses all safety related aspects of the entirety of the product and its lifecycle including both bespoke and off-the-shelf components, software, and subsystems. This specifically includes, but is not limited to:

- 1) Sensors
- 2) Actuators
- 3) Computing components

EXAMPLES: Computing hardware, operating systems, libraries

- 4) Vehicle platforms used as a basis for adding “autonomy kits”

EXAMPLE: An add-on automated driving item kit assessment requires ensuring that aspects of safety related to the underlying commercially produced vehicle have been addressed in the safety case. Credit can be taken for previous assessments of the platform, but that is likely to leave gaps unless it was specifically assessed as an EooC for the purpose of use as an autonomous vehicle platform.

- 5) On-line services

EXAMPLE: On-line map server provided from a cloud infrastructure to an operational vehicle on an as-needed basis

- 6) Logistical and maintenance support

- 7) Assumed infrastructure support

EXAMPLE: Road lane markings

e) Any aspect of a purported safety case that is not documented and/or written down, and is not provided to the assessor upon request, cannot be used in supporting a determination of conformance.

NOTE: Verbal statements and other materials not part of the safety case may be considered by the assessor to help with the assessment process, but cannot be used as supporting argument and/or evidence in evaluation of the safety case.

EXAMPLE: A developer verbally explains why a particular REQUIRED prompt element is inapplicable, but that explanation is not in the actual safety case, nor in any documentation referenced by the safety case. The lack of documented explanation for safety case deviation of that REQUIRED prompt element would cause a finding of non-conformance even though the assessor might consider the verbal explanation reasonable. The non-conformance finding could be remedied if the safety case is updated, but it is not the Assessor’s responsibility to do so.

EXAMPLE: A tester verbally states that tests were performed on a particular configuration that is being deployed. However, documentation of the configuration that was actually tested has been lost and cannot be reconstructed with reasonable certainty. To the extent that documentation of the configuration tested is necessary for a well-formed safety case, the Assessor must find non-conformance due to the missing documentation, and potentially the tests must be re-run.

f) Identify any initial portion of lifecycle during which the safety case is not intended to be valid

EXAMPLE: The safety case does not apply until the time of the first public road testing within the lifecycle.

NOTE: This permits completing the safety case in parallel with prototype development and potentially closed-course development testing. It is not intended to permit a claim of conformance for any vehicle operating on public roads at any time in its lifecycle.

g) Change management for the safety case

5.1.1.3 **HIGHLY RECOMMENDED – N/A**

5.1.1.4 **RECOMMENDED – N/A**

5.1.1.5 **CONFORMANCE:**

Conformance is checked via consideration of subsection conformance checks, including traceable inclusion of each applicable prompt element in the safety case.

5.1.1.6.1 **NOTE:** This standard was written to help support the building of a well-formed safety case. Additional measures will be required to ensure operational safety. As an example, an end product standard might require conformance not only to this standard, but to other standards involving topics such as electrical safety, fire safety, and passive occupant protection, among others.

5.1.1.6.2 **NOTE:** This clause is explicitly intended to require that the conformance documentation package is in the form of a safety case. (i.e., everything other than administrative matters dealing with the assessment itself is a part of the safety case.) As a practical matter there may be a variety of documents, repositories, and tools used to provide details such as evidence that make it impractical to have the entire safety case in a single tool or uniformly formatted data set. However, assessment requires ability to access any such material to perform an assessment.

5.1.1.6.3 **NOTE:** In general, this section places requirements upon the structure and content of the safety argument. Other major sections generally place requirements upon determining the completeness of the safety case.

REFERENCE: MISRA “Guidelines for automotive safety arguments,” 2019.

5.2 **Safety case style and format**

5.2.1 **The safety case shall use a defined, consistent format for claims, arguments, and evidence.**

5.2.1.1 **MANDATORY:**

a) Definition of claim and argument syntax, semantics, and any graphical elements used

NOTE: Notation might not have formally defined semantics. However, best available information about the safety case notation and approach should be provided to facilitate interpretation that is as uniform as practical across the developer team and assessors.

b) Definition of evidence types, formats, data dictionaries, and schemas used

c) Assessor access to complete safety case

- 1) Elements of the safety case not available to the assessor are not relied upon in conformance evaluation.

5.2.1.2 REQUIRED:

- a) Adherence to defined formats within safety case
- b) Assessor access to any available browsing, searching, reporting, and analysis tools relevant to understanding the safety case

NOTE: As a practical matter there are likely to be tools and other support used by the developers and self-auditors to make it easier to work with the safety case. Those same tools and other support are made available upon request.

- c) Identify reasoning approach regarding completeness of inductive elements of the argument

5.2.1.3 HIGHLY RECOMMENDED:

- a) Use of an established means to organize the safety case in a highly structured manner such as:
 - 1) OMG Structured Assurance Case Metamodel (SACM)
 - 2) Goal Structuring Notation (GSN)
 - 3) Claims Argument Evidence (CAE)
 - 4) Toulmin Analysis
- b) Use of tool support to aid in safety case comprehension and navigation

5.2.1.4 RECOMMENDED:

- a) Use of a graphical interface for relevant parts of safety case navigation where it increases navigability
- b) Use of structured text-based notation rather than free-form text where it increases ability to provide tool support

5.2.1.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, and design records.

5.2.1.6.1 NOTE: Regarding the word “consistent” in the clause, it is permissible for different branches of a hierarchical argument to use alternative argument formats within reason, especially if argument applies to significantly different functions or subsystems. Alternatives can make sense due to either differing original source of argument structures or the need to use a technique more suitable to specific argument technical challenges. If more than one style is used, the style applicable to each section of argument corresponds to a clearly specified style.

5.2.2 The evidence used shall conform to defined, auditable formats.

5.2.2.1 MANDATORY:

- a) A defined type for each set of evidence from a defined set of types used in the safety case

EXAMPLES: Simulation output files, test plans, vehicle data logs

NOTE: “Type” is meant in a flexible, generic sense. However, each set of evidence is of a defined type (even if each set is a unique type different from all other sets of evidence). That type is then associated with metadata that permits interpreting the evidence, per the next prompt element.

b) A defined format for each type of evidence, including at least:

1) Definition of syntax and semantics, data fields, metadata fields

NOTE: Semantics are defined to the degree practical given the type and nature of the evidence.

2) Specification of criteria for evidence consistency, correctness, and completeness suitable to make the evidence auditable

3) A defined means of validating any evidence that is not derived from acceptable data

EXAMPLE: Subjective expert judgment might be validated by data collection over the item lifecycle.

EXAMPLE: An expert opinion that lightning strikes are too rare to be relevant used as evidence for neglecting lightning strike risk mitigation is backed up by argument that any lightning strikes that do occur to deployed vehicles will be recorded and that periodic analysis will be conducted to detect if lightning strike frequency in field data becomes too frequent to neglect. (It should be noted that it is well documented that lightning does in fact strike moving, occupied vehicles upon occasion, and a more suitable safety argument is likely to be that some form of risk mitigation is in place to ensure that the post-mitigation risk from lightning strikes is acceptable.)

5.2.2.2 **REQUIRED – N/A**

5.2.2.3 **HIGHLY RECOMMENDED – N/A**

5.2.2.4 **RECOMMENDED:**

a) Descriptive or tutorial examples for interpreting each type of evidence

b) Avoidance of unconstrained free text as a data type for evidence

5.2.2.5 **CONFORMANCE:**

Conformance is checked by inspection of safety argument, evidence, and design records.

5.2.2.6.1 **NOTE:** The clause permits each piece of evidence to have a different, defined type. Using a smaller number of consistent types of evidence wherever practical can help streamline the safety case for improved comprehensibility.

5.2.3 **The claims and argument in the safety case shall be clear and consistent.**

5.2.3.1 **MANDATORY:**

a) Correct use of natural language

b) Use of a single natural language in the claims and argument

NOTE: The intent is that any natural language used in the entirety of the safety case, except for evidence, is in a single natural language appropriate for use by the item design team. Evidence can

be in alternate natural languages. A reasonable workaround for legacy multi-lingual safety cases might be the use of EooC-style component safety case interfaces between different-language sections. This does not preclude the use of additional mathematical notation and formal languages accessible to speakers of the natural language selected.

NOTE: Component safety cases do not have to be in the same language as the item level safety case. However, a component safety interface must be included in the same language as the safety case. For example, a Chinese language component safety case can export an English language component safety interface that is used in an English language item level safety case (see Section [5.7.1](#)).

c) Use of language reasonably understandable by a native speaker with general technical expertise in the item domain

NOTE: This is intended to make the safety case argument and claims accessible to an Assessor who is not part of the design team. Assessors are unlikely to be expert in the details of a particular item, but can be assumed to have general understanding of the domain and relevant technologies.

d) Avoidance of substantive ambiguity

5.2.3.2 REQUIRED:

a) Definition and consistent use of any defined notation

EXAMPLES: Formal specification language, mathematical notation

b) Use of diagrams, illustrations, interactive drawings, and other graphical approaches to supplement text when appropriate.

5.2.3.3 HIGHLY RECOMMENDED:

a) Identification of, adoption of and conformance to a technical writing style guide, including language use

b) Identification of, adoption of and conformance to a visual design language for applicable aspects of safety case related tool interfaces

c) Use of automated analysis tooling on the safety case

1) Spell checker

2) Grammar checker

3) Detection and correction of excessively complex sentences

4) Technical writing specific text style checker

5) Traceability link checkers

d) Highlighting and careful use of qualifying statements and limiting phrases

EXAMPLES: Detect and exercise care with the phrases: “essentially all,” “should be,” “generally”

e) Highlighting and careful use of negations for clarity, especially multiple negations

EXAMPLES: “Not unlike,” “A safe outcome might NOT occur”

5.2.3.4 RECOMMENDED – N/A

5.2.3.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, and design records.

5.3 Claim and argument sufficiency

5.3.1 The safety case claims shall encompass all identified safety related hazards and risks.

5.3.1.1 MANDATORY:

a) Definition of item safety requirements in the form of argument claims

1) Safety requirements for intended functionality

EXAMPLE: Avoid hitting pedestrians

2) Safety requirements for potentially unintended functionality

EXAMPLE: Vehicle movement via response to a dispatching request is unintended when in a disabled maintenance state

3) Safety requirements for unsafe behaviors and states that must be avoided

NOTE: For functional safety this is often performed via hazard analysis

4) Safety requirements for mitigating faults in the item itself, including both design faults and operational faults

5) Safety requirements for mitigating faulty, exceptional, and unspecified environmental conditions

6) Safety requirements for life cycle considerations, including updates, inspections, maintenance, and monitoring of changing operational environments

b) Mapping of each identified hazard to a potential violation of at least one relevant safety requirement

c) Identification of an acceptable level of safety for each hazard identified

5.3.1.2 REQUIRED – N/A

5.3.1.3 HIGHLY RECOMMENDED:

a) Exclusion of claims unrelated to identified hazards and risks

NOTE: This is a backward traceability requirement to avoid claims that aren't actually relevant to item safety.

NOTE: For EooC safety cases it is acceptable to trace safety case fragments to the exported EooC boundary interface under the presumption that at least some users of the EooC will have a higher level item safety case that completes the tracing relationship to identified hazards and risks.

5.3.1.4 RECOMMENDED – N/A

5.3.1.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, and design records.

5.3.1.6.1 **NOTE:** The term “claims” as used in this clause is a generic term, and includes sub-claims.

5.3.2 **The safety case argument shall support all identified claims.**

5.3.2.1 **MANDATORY:**

- a) Support safety case claims by acceptable argument
- b) Identify criteria used to determine sufficiency of arguments

5.3.2.2 **REQUIRED:**

- a) Use of epistemic defeaters in arguments

NOTE: Many of the prompt elements in other clauses are epistemic defeaters.

Additionally, reviews of each argument subtree can include consideration as to whether any additional epistemic defeaters are applicable.

- b) **Pitfall:** Taking credit for conformance to a safety standard without specifically describing the limitations of the conformance assessment is prone to over-crediting safety attributes. (*)

NOTE: This is in effect an argument gap that does not support the identified claim(s). See also Section [5.7](#).

- c) **Pitfall:** Taking credit for conformance to a safety standard designed for human operated equipment is prone to missing fault management control obligations implicitly placed upon autonomy. (*)

NOTE: This is a special, but important, case of the preceding Pitfall regarding limitations of conformance assessment to a safety standard.

EXAMPLE: Credit taken for ego vehicle controllability in assessing ISO 26262 conformance places a corresponding controllability obligation upon autonomy functions to exercise that same level of control. The need to argue a replacement for the human to provide the controllability assumed as part of ISO 26262 ASIL assignment might be missed if this Pitfall is not addressed.

5.3.2.3 **HIGHLY RECOMMENDED:**

- a) Avoid including arguments not in support of an identified claim
- b) Avoid including evidence not in support of arguments

5.3.2.4 **RECOMMENDED – N/A**

5.3.2.5 **CONFORMANCE:**

Conformance is checked by inspection of safety argument, evidence, and design records.

5.3.2.6.1 **REFERENCE:** (*) For Pitfalls see generally: Koopman, P., Kane, A. & Black, J., "Credible Autonomy Safety Arguments," Safety-Critical Systems Symposium, Bristol UK, Feb. 2019.

5.3.3 **The safety case shall avoid argument defects.**

5.3.3.1 **MANDATORY:**

UL COPYRIGHTED MATERIAL –
NOT AUTHORIZED FOR FURTHER REPRODUCTION OR
DISTRIBUTION WITHOUT PERMISSION FROM UL



Documents
Document Sharing Hub

a) Identify candidate relevant logical fallacies

NOTE: This results in a checklist of logical fallacies to be avoided in the safety case.

b) Identify candidate relevant rhetorical devices

NOTE: This results in a checklist of rhetorical devices to be avoided in the safety case.

5.3.3.2 REQUIRED:

a) Avoid identified logical fallacies

b) Avoid identified rhetorical devices

c) **Pitfall:** Taking credit for proven in use technology that is used in a different operational environment or for a different purpose is prone to over-crediting safety attributes. (*)

1) This Pitfall specifically includes COTS, legacy, and EooC components, including hardware and/or software

(**NOTE:** See Section [13.4](#))

2) This Pitfall specifically includes changes in item operational parameters that might be relevant to the component

EXAMPLE: Lions, J.L., Ariane 5 Flight 501 Failure, Report by the Inquiry Board, 1996

d) **Pitfall:** Discounting failures in field engineering feedback because there has been no previous failure is prone to inductively discounting multiple failures that, if taken as a set, substantively demonstrate invalidity of a safety case. (*)

e) **Pitfall:** Arguing coverage of autonomous failure analysis based on data from human-operated item is prone to missing some types of failures, including: (*)

1) Failures that are triggered by operational situations an autonomous item might enter that human operators typically avoid

2) Failures atypical of human mistakes that an autonomous item fault could trigger

f) **Pitfall:** Arguing risk mitigation via analysis of operational data and/or test data based on arrival rate of incidents ("surprises" or other potential failures) is prone to: (*)

1) Overlooking the additional compounding factor of the distribution of the means of different types of root causes

EXAMPLE: a heavy tail distribution of mean arrival rates of different types of triggering events for item failures

2) Overlooking the potential effects of infrequent but inevitable common cause events

EXAMPLES: Leap seconds, GPS date rollover, daylight savings time changes, or other time keeping anomalies

g) **Pitfall:** Arguing test coverage based upon human-designed test planning is prone to overlooking edge cases that apply to autonomous functions but would not generally be considered edge cases by a human operator. (*)

h) **Pitfall:** Arguing item correctness based upon use of formal methods is prone to overlooking any invalidities in underlying assumptions made by the proofs. (*)

i) **Pitfall:** Arguing that risk is low for a known hazard or variance from expected behavior based upon operational experience alone is prone to underestimating the possibility of catastrophic outcomes. Reference: Rogers Commission Report; Report of Columbia Accident Investigation Board, 1986

j) **Pitfall:** Arguing low risk based upon unvalidated simulation results alone is prone to missing risks due to simulation defects, modeling faults, and simplifications made in the abstraction process to create the simulation.

5.3.3.3 **HIGHLY RECOMMENDED – N/A**

5.3.3.4 **RECOMMENDED – N/A**

5.3.3.5 **CONFORMANCE:**

Conformance is checked by inspection of safety argument, evidence, and design records.

5.3.3.6.1 **REFERENCE:** (*) For Pitfalls see generally: Koopman, P., Kane, A. & Black, J., "Credible Autonomy Safety Arguments," Safety-Critical Systems Symposium, Bristol UK, Feb. 2019.

5.3.4 **The safety case shall avoid inclusion of defective construction patterns.**

5.3.4.1 **MANDATORY:**

a) Identify list of defective construction patterns of potential concern

NOTE: These include risky or otherwise unsuitable “patterns” in architecture, design, and implementation that are deemed unacceptable for the item. (These are not patterns actually in use, but rather a catalog of patterns that are to be avoided.)

NOTE: A minimum list includes identifying patterns enumerated in Section [5.3.4.2](#). The list can be further expanded based upon experience and literature research as determined by the design team.

5.3.4.2 **REQUIRED:**

a) Avoid identified defective construction patterns

b) **Pitfall:** Use of a “command override” pattern with a Doer/Checker architectural pattern is prone to failure to mitigate unsafe behaviors. (*)

c) **Pitfall:** Use of a Checker that does not mitigate all hazards attributable to the associated Doer is prone to failure to mitigate unsafe behaviors. (*)

5.3.4.3 **HIGHLY RECOMMENDED – N/A**

5.3.4.4 **RECOMMENDED – N/A**

5.3.4.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, and design records.

5.3.4.6.1 **REFERENCE:** (*) For Pitfalls see generally: Koopman, P., Kane, A. & Black, J., "Credible Autonomy Safety Arguments," Safety-Critical Systems Symposium, Bristol UK, Feb. 2019.

5.4 Evidence sufficiency

5.4.1 All arguments in the safety case shall be supported by evidence.

5.4.1.1 MANDATORY:

- a) Each argument element is traceable to supporting evidence.

5.4.1.2 REQUIRED – N/A

5.4.1.3 HIGHLY RECOMMENDED:

- a) Use of the following categories of evidence:

- 1) Experimental data
- 2) Analytic data
- 3) Procedure definitions
- 4) Process compliance
- 5) Development and V&V process data
- 6) V&V data

EXAMPLES: Test plans, test results, experimental design methodology, formal verification

- 7) Qualitative analysis and subjective judgement
- 8) Field engineering feedback data
- 9) Placeholder for evidence that will be collected via field engineering feedback
- 10) Accepted risks, including evidence that risk is acceptable
- 11) Assumptions for which no evidence is provided, including basis of support that the assumption is reasonable

EXAMPLE: An assumption that traffic signals never display conflicting greens might lack field data. However, a discussion with a traffic signal engineer might be cited indicating that there are hardware failsafes that disable the signal if conflicting greens are displayed, and the argument could be that this makes it reasonable to assume it will never happen in practice.

5.4.1.4 RECOMMENDED – N/A

5.4.1.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, and design records.

5.4.1.6.1 **NOTE:** Traceability can be direct or can be indirect via sub-arguments and/or sub-claims

5.4.1.6.2 **NOTE:** Argument structures may contain sub-arguments and other supporting information. However, each argument branch ultimately traces to at least one element of supporting evidence that is acceptably broad in scope to support the claims being made. Arguments that are not supported either directly or indirectly (via sub-argument) by evidence are invalid.

5.4.2 Arguments shall encompass the validity of evidence.

5.4.2.1 MANDATORY:

- a) Safety case records the experimental design or other data collection strategy for evidence based on data collection
- b) Identify criteria used to determine sufficiency of evidence
- c) Argue evidence is sufficient to result in an acceptable safety case
 - 1) Describe manner in which evidence is used to support or refute the validity of an argument and/or claim.
 - 2) Arguments that risk of confirmation bias has been mitigated

5.4.2.2 REQUIRED:

- a) Lifecycle monitoring performed upon any evidence fully or partially based upon any of:
 - 1) Unsupported expert or subjective opinion
 - 2) Existing practices that are not supported by data and are not supported by written public standards documents, public guidance documents, or similar cited sources
 - 3) Assumptions

NOTE: Lifecycle monitoring is employed to monitor risk to the degree that argument relies upon opinion, assumptions, or potentially weak evidence.
- b) Identification of epistemic defeaters and accompanying defeasibility arguments, including at least:
 - 1) Potentially confounding experimental variables
 - 2) Potential biases in data
 - 3) Potentially insufficient quantity of data samples

5.4.2.3 HIGHLY RECOMMENDED:

- a) Inclusion of counter-evidence and accompanying arguments
- b) Inclusion of process compliance arguments related to creation of evidence
- c) Lifecycle monitoring of evidence is based on consensus-based public standard approaches
- d) **Pitfall:** Collection of data before argument has been created is prone to p-hacking

5.4.2.4 RECOMMENDED:

- a) Use of adversarial sequential test design methodologies to maximize the chances of identifying failure modes and/or emergent behaviors within the ODD

5.4.2.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, evidence gathering design, and design records.

5.4.2.6.1 **NOTE:** In practice, this clause may result in an argument structure of the form: the argument is valid because (1) it is supported by evidence, and (2) the evidence itself is valid. Valid evidence is at least supported by objective, factual data.

5.4.2.6.2 **NOTE:** The potential set of epistemic defeaters and potential types of counter-evidence that might be collected is effectively unbounded. Some defeaters and counter-evidence are more likely to be relevant in practice. Selection guidance is provided by prompt elements in this standard and is additionally informed by developer experience.

5.4.2.6.3 **NOTE:** Unsupported expert opinion includes statements by domain experts that are not substantively supported by presented evidence. Opinions directly based upon scholarly papers, substantive data, and the like may be considered supported in this sense so long as that basis is stated as part of the evidence and the experimental context of the cited work is explicitly argued to apply to its use in the safety case.

5.4.3 Support of evidence validity shall encompass difficult to reproduce aspects of the item.

5.4.3.1 MANDATORY:

- a) Identification of any nondeterministic and chaotic aspects of the item with regard to evidence (if none, so state)

5.4.3.2 REQUIRED:

- a) Arguments and evidence to mitigate risk of invalidity due to nondeterministic aspects of the item and its operating environment (if none, so state)

EXAMPLE: Use of concurrency management mechanisms to mitigate timing-sensitive concurrent access faults, use of seeded pseudo-random number generators to reproduce intentionally nondeterministic behaviors for testing repeatability.

EXAMPLE: Fault injection used to fail elements which would otherwise fail infrequently during testing.

- b) Arguments and evidence to mitigate risk of invalidity due to chaotic aspects of the item and its operating environment (if none, so state)

EXAMPLE: Strategy for testing reproducibility when item behavior changes dramatically based on small input differences. For some items a concrete example is whether the vehicle veers left or right when an obstacle appears exactly in front of it.

- c) Any metrics used conform to the characteristics relating to SPI Metrics (See Section [16](#)).

5.4.3.3 HIGHLY RECOMMENDED:

- a) Use of statistical significance approaches

5.4.3.4 RECOMMENDED:

- a) Use of a digital twin or similar technique

5.4.3.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, and design records.

5.4.3.6.1 **NOTE:** Nondeterminism means that an item can behave in different ways for identical initial conditions (e.g., due to use of a pseudo-random algorithm). Chaotic means that the item can behave in different ways due to perturbations in initial conditions that are smaller than the ability of validation methods to control. Both nondeterministic and chaotic properties result in varying responses to testing, for example, but are caused by different sources. An item can be both nondeterministic and chaotic. It is not essential to differentiate between these causes if either is acknowledged to be present, but mitigation of both possibilities must be considered in establishing the validity of evidence. Non-deterministic and chaotic behavior can undermine the validity and completeness of test results if tests pass by chance due to favorable system behavior in a particular test run.

See also Section [12.4.7](#) regarding fault injection.

5.5 Accepted risks

5.5.1 Accepted risks shall be identified.

5.5.1.1 MANDATORY:

- a) Identify acceptance criteria for any risk less than fully mitigated.

5.5.1.2 REQUIRED:

- a) Identify any risk that is less than fully mitigated to an acceptable level as an “accepted risk.” These include but are not limited to:

- 1) Risks for which no mitigation measure has been taken

EXAMPLE: An item level safety assessment potential hazard that was determined to be extremely improbable or otherwise not something that could happen in the “real world” is an accepted risk and is included in the safety case as an unmitigated risk.

- 2) Partially mitigated risks (i.e., risks not mitigated to a fully acceptable level)

- 3) Unknown risks

EXAMPLES: “known unknowns,” “unknown unknowns”

- 4) Risks for which mitigation is based on unsupported expert opinion, assumptions, or other less-than-comprehensive evidence and also are not in conformance with generally accepted industry practices such as standards documents

- b) For less than fully mitigated risks (including accepted risks) characterize the level of mitigation and argue that the level of mitigation, if any, is acceptable in the context of the item safety case.

- 1) Include an evaluation of the expected outcomes of each such risk across the item lifecycle.

5.5.1.3 HIGHLY RECOMMENDED:

- a) Characterization of post-mitigation level of risks that are deemed fully mitigated.

5.5.1.4 **RECOMMENDED – N/A**

5.5.1.5 **CONFORMANCE:**

Conformance is checked by inspection of safety argument, evidence, and design records.

5.5.1.6.1 **NOTE:** This includes acceptance of any remaining unknown risks, which should be broken down into categories if and as information exists to support doing so (e.g., “known unknowns”). It is understood that such risks might not be readily quantified, but yet they are being accepted when a decision to deploy is being made.

5.5.2 **Accepted risks shall be tracked through the item lifecycle via field engineering feedback**

5.5.2.1 **MANDATORY – N/A**

5.5.2.2 **REQUIRED:**

- a) Field engineering feedback for accepted risks to ensure that the risk in practice is less than or equal to the level of risk stated as an expected outcome for item operation
- b) Field engineering feedback explicitly considers risks due to gaps in risk analysis

5.5.2.3 **HIGHLY RECOMMENDED:**

- a) Automated field data collection of incidents and loss events to determine if outcomes for accepted risks are within expectations on a per-risk basis

5.5.2.4 **RECOMMENDED – N/A**

5.5.2.5 **CONFORMANCE:**

Conformance is checked by inspection of safety argument, evidence, and design records.

5.5.2.6.1 **NOTE:** An example approach is including “unknown risks” in the safety case, noting them as “accepted risks,” justifying why acceptable risk analysis has been performed to permit accepting the resultant risks, and arguing that the root cause analysis procedure explicitly includes identifying novel risks so as to add them to the safety case in a timely fashion when they are encountered in the field.

5.6 **Safety culture**

5.6.1 **The role of safety culture of development, supply chain, maintenance and operations in risk identification and mitigation shall be identified.**

5.6.1.1 **MANDATORY:**

- a) Definition of safety culture and role used as part of the risk mitigation approach, addressing at least:
 - 1) Role of safety staff in ensuring safety
 - 2) Role of non-safety staff in ensuring safety
 - 3) Role of management in ensuring safety
 - 4) Role of safety management system in organization

- 5) Role of suppliers in ensuring safety
- 6) Role of lifecycle participants in ensuring safety
- 7) Independence of safety roles between engineering development stakeholders, deployment stakeholders, and business profitability stakeholders
- 8) Upper management visibility of and delegation of authority to safety roles

b) Identify activities that support the communication and tracking to resolution of potentially safety related issues

c) Identify an acceptable set of ongoing field engineering feedback and continuous improvement activities related to safety culture

d) Identify an acceptable set of ongoing activities to gather information on hazards and risks from publicly available sources

EXAMPLES: Monitoring recall notices from other developers, accident investigation reports, regulatory actions in the system domain plus related domains, published statistical analyses of hazards.

e) Argue that the execution of identified activities and other identified factors results in an acceptable safety culture

5.6.1.2 REQUIRED:

a) **Pitfall:** Organizational structures in which engineering, business, and/or operational management can exert control or pressure upon roles tasked with ensuring safety are prone to degraded safety outcomes.

EXAMPLE: Normalization of deviance

b) All activities that have access to safety related data identified as supporting the communication of potentially safety related issues

c) Identification of role in safety culture in ensuring that root cause analysis will be effective at identifying defects in safety cases and safety processes

5.6.1.3 HIGHLY RECOMMENDED:

a) Identification of metrics and feedback mechanisms used to evaluate and manage safety culture

b) Identification of roles and responsibilities accompanied by argument and evidence for suitable competency of staff

EXAMPLE: Evidence supporting an argument of Suitably Qualified and Experienced Personnel (SQEP)

5.6.1.4 RECOMMENDED – N/A

5.6.1.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, and design records.

REFERENCE: See safety culture metrics for SPIs in Section [16.2.5](#)

5.7 Item scope

5.7.1 The argument shall identify safety related aspects of the item, including potential faults and failures, encompassing the item lifecycle.

5.7.1.1 MANDATORY:

- a) Identify and describe safety related functionality
- b) Identify and describe safety related components
- c) Identify and describe safety related properties

EXAMPLE: Real time deadline for responding to hazards when credit is taken for response time in risk reduction argument

- d) Identify and describe safety related aspects of non-operational lifecycle phases (See Section [14](#))
- e) Identify and describe other aspects of the item related to safety

EXAMPLE: Emergent properties such as total system weight if limited kinetic energy is part of the risk mitigation strategy

- f) Functional requirements for safety related functionality
- g) Acceptable coverage of identified fault models (See Section [6.2](#))
- h) Coverage analysis for the V&V of each identified element

NOTE: See subsections within [12.3](#) for more information

5.7.1.2 REQUIRED:

- a) Safety related exception handling capabilities
- b) **Pitfall:** Components and functions that provide data to or otherwise affect the operation of safety related components and functions are themselves safety related, but are prone to being discounted as not safety related if traceability of data flows and other direct and indirect sources of interaction with safety related components is not performed rigorously.

5.7.1.3 HIGHLY RECOMMENDED – N/A

5.7.1.4 RECOMMENDED – N/A

5.7.1.5 CONFORMANCE:

Conformance is checked via consideration of subsection conformance checks, including traceable inclusion of each applicable prompt element in the safety case.

5.7.1.6.1 NOTE: A primary goal of this clause is to ensure that aspects of the item that can contribute to or mitigate risk are identified to ensure inclusion in the safety case.

5.7.2 The safety case shall describe the concept of operations for the item.

5.7.2.1 MANDATORY:

- a) Overview of mission capabilities
- b) Overview of safety objectives
- c) Overview of risk mitigation strategy
- d) Overview of item hardware, software, and functional architecture
- e) Overview of safety related functionality
- f) Overview of item operational modes

5.7.2.2 **REQUIRED – N/A**

5.7.2.3 **HIGHLY RECOMMENDED:**

- a) Overview of non-safety related functionality and overview of non-interference explanation with safety related functionality

5.7.2.4 **RECOMMENDED – N/A**

5.7.2.5 **CONFORMANCE:**

Conformance is checked via inspection of the safety case.

5.7.2.6.1 **NOTE:** This clause is intended to provide an overview of the item and specifically to provide contextual information needed for understanding the safety case.

5.7.3 The boundary within the safety case between any assessed Element out of Context (EooC) and the rest of the safety case shall include a specified interface.

5.7.3.1 **MANDATORY – N/A**

5.7.3.2 **REQUIRED:**

- a) Identification of EooC boundaries in the safety case, if any
- b) For each EooC boundary:
 - 1) List of EooC properties that have been assessed
 - 2) List of EooC assumptions that must be true for assessed properties to hold true
 - 3) EooC fault model
 - 4) Identify any UL 4600 clauses that were not considered in EooC assessment.

EXAMPLE: A purely hardware EooC might consider software-related clauses out of scope for EooC assessment. This would put the full burden of conformance with those clauses onto the safety case incorporating the EooC.

- 5) Assessment report for EooC in same language as item safety case using the assessment report

NOTE: This might require a translation of the EooC safety report from the language used for the component safety case to the language used in the item safety case

6) All other characteristics of the EooC that are used as evidence by the safety case for the item that contains the EooC.

5.7.3.3 HIGHLY RECOMMENDED – N/A

5.7.3.4 RECOMMENDED – N/A

5.7.3.5 CONFORMANCE:

Conformance is checked by inspection of safety argument, evidence, and design records.

5.7.3.6.1 **NOTE:** The concept of EooC used here is generic and not limited to the functional safety aspects of the component as might be the case when the term Safety Element out of Context (SEooC) is used in a functional safety standard. Assessment of an EooC is performed for the full scope of its use in the safety argument.

5.7.3.6.2 **NOTE:** Lists of EooC assumptions and properties are documented to the degree they are known. As additional assumptions and properties are discovered over the component lifecycle, they are added. Addition of an additional assumption or property might invalidate higher level safety cases using the EooC. **Reference:** Lions, J., Ariane 5 Flight 501 Failure Report by the Inquiry Board, 19 July 1996.

5.7.3.6.3 **NOTE:** Other approaches to modular safety argument, re-use of argument fragments, inter-domain argument reuse, argument patterns, and so on are not specifically prohibited. However, establishing assessment criteria for alternate approaches is beyond the scope of this standard.

6 Risk Assessment

6.1 General

6.1.1 The safety case shall identify risks and argue acceptable mitigation.

6.1.1.1 MANDATORY:

- a) Identification of fault models (See Section [6.2](#))
- b) Identification of hazards (See Section [6.3](#))
- c) Risk evaluation (See Section [6.4](#))
- d) Risk mitigation and evaluation of mitigation effectiveness (See Section [6.5](#))

6.1.1.2 REQUIRED – N/A

6.1.1.3 HIGHLY RECOMMENDED:

- a) Identification and use of a total item risk summing approach
- b) Calculated probability of incident occurrence yields both probability and confidence for entire causal chain. Probability is considered unbounded unless coupled to confidence.

6.1.1.4 RECOMMENDED:

- a) Bayesian probability estimation for risk approaches that require estimating probability of occurrence

6.1.1.5 CONFORMANCE:

Conformance is checked via consideration of subsection conformance checks, including traceable inclusion of each applicable prompt element in the safety case.

6.1.1.6.1 NOTE: Risks that are accepted are still monitored as a mitigation for a situation in which the acceptance was based on an assumption that is incorrect or becomes incorrect. Thus, all risks are said to have been “mitigated” in the final safety case even if there is no explicit mitigation mechanism included in the item design.

See also: Run-Time Monitoring, Section [12.5](#).

6.2 Fault model**6.2.1 The argument shall define a fault model for safety related aspects of the item.****6.2.1.1 MANDATORY:**

- a) Fault model identified for each safety related item component, feature, or other aspect for which fault analysis is relevant
- b) Design fault model
 - 1) Hardware fault model, including microelectronics
 - 2) Software fault model
- c) Manufacturing fault model
- d) Operational fault model
- e) Non-operational fault model

EXAMPLES: Age-related component degradation, degradation of item due to lack of operation while in storage

- f) Maintenance fault model
- g) Procedural fault model
- h) Item operation fault model
- i) Tool fault model
- j) Random faults

k) Systematic faults

l) Fault multiplicity

1) Single fault

2) Multiple faults due to a common cause

3) Accumulation of multiple faults over the lifetime of the item

NOTE: Credit can be taken for diagnosis, recovery, degraded operational modes, and repair capabilities if supported by evidence

m) Undetected (latent) faults

n) Permanent, transient, and intermittent faults

o) Traceability from each fault model to fault mitigation for applicable components, functions, and other aspects of the item.

6.2.1.2 REQUIRED:

a) Model of safety related expectations for correct or required operation for each safety related component or function.

NOTE: In some safety cases this might correspond to “safety requirements,” but it is potentially broader if for example there are safety related non-functional aspects that need to be considered.

EXAMPLE: Abnormally excessive power consumption by a fail-safe device depletes a battery power supply, disabling a failsafe. In this case specified power consumption is a safety related expectation that might not normally be associated with a “safety requirement.”

b) Fail arbitrary fault model at the component level for complex electronic components within an FCR.

EXAMPLES: Programmable hardware, components containing software

NOTE: An arbitrary failure mode at the component level is an arbitrary fault symptom at the next higher level of abstraction, and therefore is being referred to as an arbitrary fault in this context.

c) Inclusion of “other” fault modes for non-complex electronic component fault analysis

EXAMPLE: Beyond “short” and “open” to include resistive faults and capacitive faults for connectors and passive electronic components

d) **Pitfall:** Simplistic fault models are prone to being unacceptable for describing faults in computer-based items

EXAMPLES: Assumptions that all component failures result in clean fail-stop semantics, considering only input shorts to power/ground, considering only fail-stuck integrated circuit faults, and other fault models that are overly simplistic for some applications

6.2.1.3 HIGHLY RECOMMENDED:

a) Inclusion of electrical, mechanical, and other components that are relevant to ability of autonomy to operate safely

EXAMPLES: Sensor mounting structure, vehicle wheels

- b) Encompassing detected but un-annunciated faults
- c) Use of Byzantine component fault model (includes hazardously misleading information)
- d) **Pitfall:** Fault models limited by conclusory and subjective statements that a typical type of fault is “unrealistic” or would “not happen in the real world” are prone to significantly understating faults that actually do happen in fielded items.

EXAMPLE: Exclusion of Byzantine faults in aerospace item fault models (see Driscoll, K., “Real Item Failures,” <https://c3.nasa.gov/dashlink/resources/624/>)

6.2.1.4 **RECOMMENDED – N/A**

6.2.1.5 **CONFORMANCE:**

Conformance is checked via consideration of subsection conformance checks, including traceable inclusion of each applicable prompt element in the safety case.

See also: Fault Detection and Mitigation, Section [10.4](#)

See also: Malicious Fault Model, Section [10.8.2](#)

6.2.2 **The software fault model shall include an acceptably broad set of potential software faults and failures.**

6.2.2.1 **MANDATORY:**

- a) Tool chain failure

EXAMPLE: Tool produces incorrect software image or configuration

- b) Incorrect requirements and algorithms

- c) Incorrectly built software image including nonvolatile data

EXAMPLE: Wrong version of library included

- d) Incorrect installed software image, including nonvolatile data

EXAMPLE: Deployed software image differs from validated software image

- e) Incorrect software image version information, including nonvolatile data

- f) Corrupted software image, including nonvolatile data

- g) Data reporting tool failure

EXAMPLE: Incident data reporting tool reports incorrectly

- h) Coding defects

- i) Defects that corrupt data at run time

- j) Defects that corrupt hardware runtime configuration

- k) Timing faults

- l) Race condition

6.2.2.2 REQUIRED:

- a) Faults resulting from use of hardware description languages (HDLs)
- b) Third party component failure

EXAMPLE: RTOS defect, math library defect

- c) Configuration reporting tool failure
- d) Tool calibration and usage errors

EXAMPLES: Incorrect calibration variable definitions used with configuration data compilation tool, incorrect command line parameters used when running static analysis tool or compiler

- e) Faults and failures associated with entry into and execution of supervisory sections of safety related software

EXAMPLES: Concurrency faults, timing faults, and data concurrency faults that might be caused by interrupt service routines and/or operating system scheduling; disabling interrupts in a way that disrupts system scheduling

- f) Defects in software embedded in third party components

EXAMPLES: Sensor firmware defects, network adapter firmware defects, storage module firmware defects

- g) Machine learning brittleness
- h) Corrupted data
- i) Boundary value faults

EXAMPLES: Off-by-one, incorrect handling of borderline cases, array index values at boundaries of array size

- j) Exceptional value faults

EXAMPLES: Null pointers, floating point infinity, zero length character strings

- k) Out of range value faults

EXAMPLES: Off-scale sensor values, buffer overflows, numeric overflow, floating point underflow

- l) Software faults that corrupt memory

- 1) Data memory

- i) Volatile memory
 - ii) Non-volatile memory
 - iii) Stack
 - iv) Heap
 - v) Statically allocated data

- 2) Program memory

m) Software faults that result in incorrect hardware performance

EXAMPLES: Software corruption of hardware configuration, including clock rate dividers, I/O pin configurations, and power management settings

n) Software defects that defeat integrity checks

EXAMPLE: Spurious watchdog timer kick

o) Software component crash

p) Software component hang

q) Software component misses real time deadline

r) Time keeping faults, including:

1) Time keeping overflow

EXAMPLE: Timer rollover

2) Time keeping roundoff

EXAMPLE: 32-bit floating point representation of time suffers roundoff error when incremented by tenths of a second

3) Incorrect handling of time keeping discontinuities

EXAMPLES: Leap second, leap year, "Y2K" bug, Unix time rollover in 2038, Daylight Savings Time, GPS week rollover

4) Incorrect handling of time zones

EXAMPLES: Crossing international date line, crossing time zone boundaries, changes to time zone boundaries and/or offsets

5) Incorrect calculation of local solar time and position

EXAMPLE: ODD precludes driving with the sun visible on the horizon due to potential camera issues, but item's calculation of apparent sun angle at the item's current position (latitude, longitude, altitude) and date is incorrect due to incorrect calculations and/or incorrect nautical almanac stored data values.

s) **Pitfall:** Use of a fail-crash software fault model is prone to overlooking fail active and other dangerous software failure behaviors.

6.2.2.3 HIGHLY RECOMMENDED:

a) Other sources of systematic faults and errors

b) A specified fault model for malicious data faults

c) An unbounded model for malicious data faults with the sole exception of adversary not knowing secret key values

6.2.2.4 RECOMMENDED – N/A

6.2.2.5 CONFORMANCE:

Conformance is checked via inspection of design and V&V evidence as well as demonstration.

6.2.3 The microelectronic and electronic hardware fault model shall cover an acceptably broad set of potential run-time as well as fabrication faults and failures.

6.2.3.1 MANDATORY:

a) Power supply faults

b) Thermal faults

EXAMPLES: Clock throttling due to excessive temperature, early failure due to chronic operation at or above maximum rated temperature

c) IC fabrication faults

d) Embedded firmware defects

e) Single event upsets

6.2.3.2 REQUIRED:

a) Clocking faults

b) Errata and design defects

c) Specific types of power supply faults:

1) Under-voltage (brownout)

EXAMPLES: Incorrectly configured brown-out protection; brown-out threshold voltage set for microcontroller requirements but is too low for associated non-volatile memory chip causing memory corruption when writing

2) Fast transient power spikes that do not activate brownout protection

3) Over-voltage

i) Due to regulator failure

ii) Due to external application of high voltage

EXAMPLE: Tow truck applies 24V jump voltage to 12V electrical system

iii) Due to power cabling faults

4) Power phasing and polarity faults

i) Reverse polarity battery installation

ii) Application of DC external voltage with reverse polarity

iii) Coupling with mismatch AC phase angle

iv) Coupling with incorrect AC 3-phase rotation direction

5) Loss of at least one supply voltage