Rationalising safety cases for naval systems

J R Inge[1*] CEng CITP MIET MBCS MAPM, D Gardner[1] CEng MIET, C Brooking[2] MEng CEng FIMechE CEnv AIEMA FSP

[1]*Defence Equipment & Support, UK*
[2]*Occam Group Ltd, UK*
*Corresponding author. Email: james.inge@scsc.uk

Synopsis

As naval systems become more complex, it is increasingly challenging to provide assurance that they can be operated safely. Safety cases need to be cost-effective to produce, yet robust in delivering a well-founded argument for the safety of the overall capability system. There is the potential to spend disproportionate effort demonstrating the safety of relatively simple, well-understood equipment; while not necessarily applying enough effort to understand how system elements function together to deliver a safe overall system. Increasingly, naval capabilities are assembled as a system-of-systems, bringing together a mix of bespoke, off-the-shelf and legacy elements, including both onboard and offboard systems. Such complex systems need a systems engineering approach to system safety. This paper examines some of the work underway to help rationalise and streamline management of safety cases for complex systems, including IEC 63187, the international standard currently being drafted for systems engineering, system safety and complex systems in defence applications; and Def Stan 02-904, the new UK Defence Standard on Surface Ship Safety Critical Items.

Keywords: Systems engineering; System safety; Safety-critical items; Standardization; Safety cases

1. Introduction

1.1. The problem with safety cases

To gain assurance that the systems it procures will be safe, the UK Ministry of Defence (MOD) has for many years required the production of a safety case (Inge, 2007). A safety case is defined as 'a structured argument, supported by a body of evidence that provides a compelling, comprehensible and valid case that a product, service or system is safe for a given application in a given environment' (MOD, 2023a). This argument changes through the system life cycle: starting as a case that it will be feasible to make a capability safe, then that the product delivering the capability is designed and manufactured in a way that will make it safe to operate, then that it is maintained in a state where this is so, and actually operated safely. A safety case is documented in a series of safety case reports, which summarise the arguments and evidence at a moment in time and document progress towards implementing the safety strategy. However, there is considerable flexibility on exactly what is required to constitute an acceptable safety case, in terms of its format, structure, and level of rigour. This flexibility allows an appropriate approach to be tailored to each project but introduces an ambiguity about the requirement that can lead to inefficiency.

Military systems are often very complex, both in a technological and a managerial sense. Naval platforms are often physically large, with a multitude of subsystems that can themselves be complex engineered artefacts, such as weapon systems or sensor suites. While some of these subsystems may be developed as part of the platform design or in parallel projects, in many cases they are developed separately. Some may be commodity parts or off-the-shelf designs that are already in-service elsewhere, others may be legacy equipment cross-decked from other platforms. And with ships in service for decades, they are likely to be fitted with new subsystems in the future, which may not have been considered at the time they were built.

Authors' Biographies

James Inge leads the Ships domain Safety and Environmental Protection Team in Defence Equipment & Support. He is a past chair of the MOD Safety and Environmental Standards Review Committee and is currently part of IEC SC65A Working Group 18, developing the new system safety standard for defence, IEC 63187.

Dan Gardner is Deputy Head of Engineering and Chief Marine Electrical Engineer in the Ships Engineering HQ at DE&S. He has a background in the offshore, nuclear and renewables industries, and has been leading the work to develop Def Stan 02-904.

Charles Brooking is a principal consultant at Occam Group Ltd, providing safety assurance services for complex systems in the defence maritime, weapons, land and nuclear domains.

 https://doi.org/10.24868/11189

All too often though, safety cases are written as a set of bolt-on documentation for an existing product. In many cases, the authors are consultants who do not work directly for the Original Equipment Manufacturer (OEM) or integrator. There may be good reasons for this: suppliers may not be familiar with the safety case regime; some manufacturers may no longer exist for legacy systems. However, this can mean the authors are given little access to information about the design or the context of use and the safety case report becomes detached from the engineering processes or the operation of the system; it may reflect neither the good safety engineering that might have been done by the designer, nor the true hazards seen in operation. Without having clear links between safety arguments at different levels of the systems hierarchy, it can be hard to understand what attributes of a sub-system make it safe – which can make it hard to replace if an alternative with identical form, fit and function is not available. Conversely, we may be unaware of latent safety problems, if it is not clear what role a sub-system was supposed to play in the safety case of the higher-level system.

In the past, the safety case approach has sometimes been criticised as producing 'impenetrable tomes' (SDF, 2022), or obscure documents of bureaucratic length (Haddon-Cave, 2009). One frequently finds lengthy safety case reports for minor equipment that pose very little hazard, while all too often, accident investigations for more complex systems find issues that had been overlooked in the safety case. To paraphrase Tony Hoare, it is difficult to make a safety case that is so simple that there are *obviously* no deficiencies; it is far easier to make it so complicated that there are no *obvious* deficiencies (Hoare, 1981). There is a risk that a considerable amount of time, effort and expense is spent creating and maintaining safety cases that do not do much to improve safety.

The challenge is to avoid on the one hand producing reams of paperwork that add little value, and on the other hand giving superficial treatment to important safety issues. In the past, there was a perception that every item of equipment the MOD purchased required its own exhaustive safety case. The current Defence Maritime Regulations, DSA02-DMR, now recommend that the safety case is 'developed proportionate to the perceived level of safety risk' (MOD, 2024). Producing a 'proportionate' safety case means finding the happy medium between the extremes of being recklessly scant and paralyzingly comprehensive. But how should this best be achieved?

1.2. The structured argument approach

A common approach to producing a safety case is to use a structured argument. In this context, a structured argument is one in which makes a high-level claim and breaks it down into a hierarchy of sub-arguments that are eventually supported by evidence. This may be presented as structured text, or graphically to make it easier to understand the relationships between elements of the argument. Two graphical notations often used for this purpose are the Goal Structuring Notation (GSN) (ACWG, 2021), and Claims, Argument and Evidence (CAE) (Adelard, 2024). Typically, the top-level argument is a statement such as 'the [system] is acceptably safe for its [context of use]', with appropriate explanations to describe the scope of what is meant by the [system] and the [context of use]. In GSN, this is known as a 'goal', since the goal of the safety effort is to ensure that the statement is true. In CAE, it is a 'claim', which expresses the conclusion that the rest of the argument is expected to support. The goal/claim is then broken down into sub-goals/claims, until the argument is defined clearly enough that the lowest-level goals can be satisfied by tangible pieces of evidence ('solutions' in GSN terms). At a basic level, GSN and CAE are similar, but GSN includes additional elements to help explain the context, assumptions, justifications and strategies underpinning complex arguments.

GSN also includes features for modularising safety arguments, which can be helpful for complex systems. MOD policy in Joint Service Publication JSP 815 (MOD, 2023b) and DSA02-DMR requires the safety case to be endorsed by a representative of the operator: either the Senior Responsible Owner (SRO) of the programme acquiring a system, or the Duty Holder or other Accountable Person responsible for operating it once in service. For a complex system such as a warship, it is not practical to expect the operator to sign off separate safety cases for each component system element. Component-level safety cases need to be integrated into an overall safety argument for the system. Modular safety cases in GSN provide a mechanism to do this: an 'away goal' in a platform safety case may be supported by the argument presented in a separate safety case module, e.g. the safety case for a sub-system.

Safety is an emergent property resulting from interactions between all the elements of a complex system. While the principle of 'platform primacy' suggests that the safety case for a naval platform should incorporate arguments for the safety of its component elements, it is not sufficient just to claim that the platform is safe because its components are safe. We also need to argue about the way that sub-systems are integrated: how they work together to deliver safety functions, and how they avoid interfering with each other in a hazardous manner. This means that there will normally need to be multiple touchpoints between a platform-level safety argument and its supporting sub-system safety arguments.

 https://doi.org/10.24868/11189

GSN allows interfaces between modules to be defined by 'contracts'. These describe the links between the away goals that require support, their relevant context and justifications, and the goals, context and justifications in other module(s) that provide the supporting argument (ACWG, 2021). As well as making it clear how different safety argument modules support each other, formally describing these interfaces can make it easier to swap modules, e.g. by facilitating analysis of whether the safety case for a new sub-system still supports the safety case for the platform in the same way as its predecessor.

Structured arguments presented using notations such as GSN and CAE can be useful to help clearly set out the safety argument. At the start of a project they make it easier to see what evidence will need to be generated to argue that a PSS is safe; later they can help identify the impact if the system changes or the validity of evidence is brought into question. And breaking structured arguments into modules can help manage complexity by abstracting the detailed arguments about subsystems. However, to be effective and efficient, we must choose a good argument structure that is proportionate to the risk involved.

2. Efficient arguments for maritime safety

2.1. Choosing the argument structure

Depending on the stage of the project and the purpose of the safety case report, it may be appropriate to break the top-level argument for safety of a product, service or system into two sub-goals: that the PSS is 'safe to operate', and that it is 'operated safely'. Within the MOD, this neatly breaks the argument into the parts managed by the acquisition organisation (normally Defence Equipment & Support or the Submarine Delivery Agency), and the parts managed by the system operator, for instance the Royal Navy.

Below the top level of the safety argument, one needs to choose an appropriate strategy to make a compelling argument that the system is safe. One strategy is to argue about the effectiveness of the Risk Control Systems (RCS) that control the risk posed by the product, service or system as part of a Safety and Environmental Management System (SEMS). In DSA02-DMR, the Defence Maritime Regulator recommends adopting appropriate risk control systems from Table 1.

Table 1. DSA02-DMR Risk Control Systems (MOD, 2024).

RCS	Description
a)	Safety and Environmental Management System documented using Organisation and Arrangements Statements and Safety and Environmental Management Plans
b)	Certification and Certification Strategy
c)	Integration of Safe Design and Construction
d)	Maintenance of the Ship and Equipment
e)	Management of Change; Maintenance of Conditions
f)	Documentation
g)	Crewing Levels, Competence and Training
h)	Incident Reporting and Analysis
i)	Emergency Preparedness
j)	Safe and Environmentally Compliant Operating Envelope
k)	Live Health, Safety and/or Environmental Case and the Health, Safety and/or Environmental Case Report, Summaries and Statements
l)	Requirements Management
m)	Verification of Internal Assurance (1st Party Assurance and 2nd Party Assurance)

Unfortunately, such an approach can be inefficient, especially for relatively simple pieces of equipment. The risk control systems used in DSA02-DMR are based on work that was originally intended to give assurance about the end-to-end safety process across organisations within a complex enterprise (Inge and Costello, 2008). With sufficient interpretation (DSA02-DMR does not expand beyond the descriptions in Table 1), each risk control system can be relevant to a safety case, but the set is not optimised for this purpose. This can lead to a temptation to write a similar amount for each one, rather than focus on those that are most important, or to fill in the space even where a theme is not relevant. The set of risk control systems in DSA02-DMR acts more as a list of relevant topics than a structure for a compelling argument.

 https://doi.org/10.24868/11189

An alternative argument strategy, illustrated in GSN at Figure 1, has been used for some while for Type Airworthiness Safety Assessments (TASA) in the Air domain in DE&S and can also be adapted for the Maritime environment. The top-level goal that a Product, Service or System (PSS) is acceptably safe, G0, is broken into four sub-goals. Guidance in the DE&S Air Engineers' Toolkit breaks these sub-goals down to a further level of detail, with the requirements varying between commodities and more complex systems.

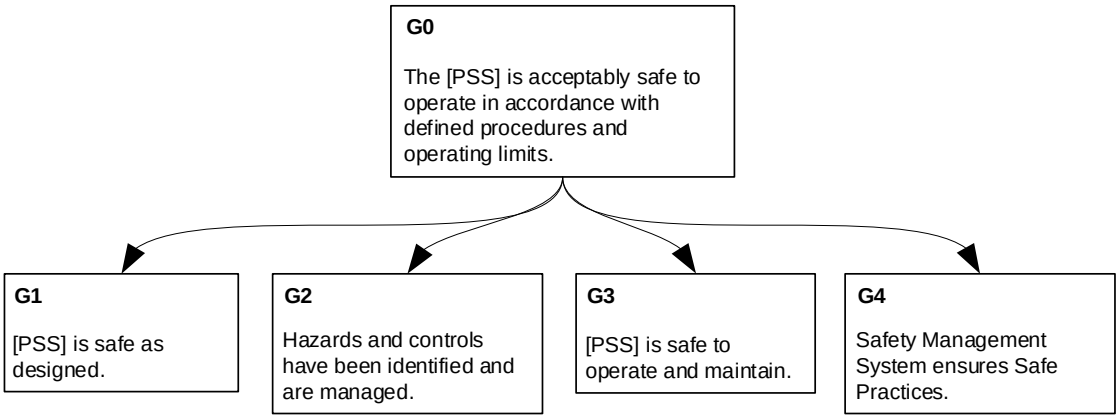

Figure 1. Four-pillar safety argument, based on Air Engineers' Toolkit TASA model.

Under G1, arguments are made that the designer or supplier has conducted a suitable safety process themself, certification standards are met, the design has been tested in a representative environment for its intended use, and is kept under proper configuration control. G2 argues that an effective risk analysis process has been put into effect. G3 makes the claim that the user has been provided with the information necessary to safely use the product, service or system (via documentation, training, emergency arrangements, etc. as appropriate). G4 makes a process-based argument that management processes are in place to review and update the safety case, and to ensure compliance with relevant standards, legislation, regulation and policy. The goals in Figure 1 overlap the topics covered by the risk control systems at Table 1, but do not correspond directly to them. Some risk control systems, such as 'Documentation', impact multiple goals, while others like 'Certification and Certification Strategy' are more specific to one area of the argument.

Beneath the four sub-goals, one must decide how to expand the argument in a way that is proportionate to the circumstances in question. This must consider the complexity of the product, service or system under consideration and its interfaces with the wider system it operates in, as well as the magnitude of the hazard and consequent need for reassurance about safety. The argument must also be tailored to the stage of the system in its product life cycle. Early in the acquisition phase, the argument will focus more on G1, checking that it is designed or selected appropriately. As the system comes into service, the argument shifts more to G2 and G3, checking that safety arrangements are in place and are operating effectively.

2.2. Arguments for low risk / low complexity equipment

JSP 815 guides that for some equipment, a safety case approach may not be proportionate, if the complexity and level of risk involved is sufficiently low (MOD, 2023b). This might apply to equipment where no specific hazard or contribution to safety is identified. Some kind of safety assessment is still required, but it can be argued it would be proportionate to do this at the level of groups of related equipment, rather than individual items, and to base the assessment more on the management processes involved than technical analysis of a Product, Service or System (PSS). Such a safety argument could be developed along the lines of Table 2 (shown in tabular format for brevity – graphical notation could also be used).

 https://doi.org/10.24868/11189

Table 2. Skeleton safety argument for low risk/low complexity items with negligible safety impact.

Goal	Supporting lines of argument	Examples of potential supporting evidence
G1	The appropriate standards for the PSS are well understood, and there is evidence of compliance.	CE/UKCA marking, Original Equipment Manufacturer (OEM) endorsement of spare part.
G2	The PSS is not known to be safety-related, and no specific hazards have been identified.	Declaration by Suitably Qualified and Experienced Personnel (SQEP) holding appropriate safety delegation.
G3	There is no requirement for specific operating or maintenance information to ensure safety.	
G4	Systems are in place to select competent suppliers, manage product quality, review ongoing suitability.	Safety Management System documentation. In-service defect or incident reports.

The bulk of a safety case based around Table 2 would be focused on Goal 4, arguing that the acquisition organisation and operators had appropriate processes in place. The only item-specific parts of the case would be the list of items to which it applied and the records that the SQEP delegated person had assessed that this style of argument was indeed appropriate to those items. This approach of using the skeleton argument from Table 2 to create a generic safety case applicable to a list of items would only be appropriate for the lowest-risk and complexity items. Once items are known to have a safety impact, more specific information is required to justify their safety, as shown in Table 3.

Table 3. Skeleton safety argument for low risk/low complexity items.

Goal	Supporting lines of argument	Examples of potential supporting evidence
G1	Appropriate standards for the PSS have been selected and complied with.	Declaration by SQEP delegation-holder that chosen standards are appropriate for the application. Manufacturers' declaration of conformance. Type approval certificates. Test/survey results and certification. List of hazardous materials.
G2	All identified hazards have been controlled via adherence to the appropriate standards, or supply of appropriate safety information.	Declaration by SQEP delegation-holder.
G3	Appropriate safety information is included in training, manuals, operating limitations, etc.	References to the relevant documents.
G4	Systems are in place to select competent suppliers, manage product quality, review ongoing suitability.	Safety Management System documentation. In-service defect or incident reports. Audit reports.

The skeleton argument outlined at Table 3 may be appropriate where the product, service or system is known to have a safety impact, but that impact is well controlled by existing standards and processes. It is not sufficient to use a generic safety argument as described previously, but it may be proportionate to apply a template argument to groups of similar systems and populate it with references to specific evidence for each item. This saves the effort of generating a new argument from scratch for each item. To further reduce duplicated effort, the Safety Management System argument at G4 can be referred to as a safety argument module and re-used between groups of different system types.

There is a risk here: the safety case approach has been criticised for a tendency towards confirmation bias, or assuming that the top-level goal is true without critically analysing the safety argument (Haddon-Cave 2009; Leveson, 2011). When using a template safety case, it is vital to populate it with details that are relevant to the system in question and are verified to be true. For instance, it is not sufficient to just say that a hazard is controlled by 'documentation'; it needs to be clear which documents are being referred to, and that those documents actually hold the relevant details for that system. Similarly, when basing a safety argument on the existence of a documented Safety Management System, there needs to be evidence that the procedures in that management system were actually applied to the product in question.

 https://doi.org/10.24868/11189

2.3. Safety Critical Items approach

As system elements become more complex and have more interactions with higher-level and peer systems, it becomes harder to support the goals that 'the system is safe as designed' and the 'hazards and controls are identified and mitigated' (G1 and G2 in Figure 1). Platform and system-level concerns often start to dominate, and safety becomes more dependent on integration aspects rather than individual system element designs. Addressing these integration issues can be a management challenge, and if there is insufficient delineation between local hazards to personnel and hazards with platform-level effects, it is easy to focus effort on the wrong issues.

With *Def Stan 02-904 – Surface Ships Safety Critical Items* (MOD, 2023c), the MOD is starting to adopt an approach for naval ships that is already in use for submarines and in the offshore oil and gas industry. For certain system elements that are designated as 'Safety Critical Items', i.e. those that are essential to the platform 'safe to operate' argument, Def Stan 02-904 requires the associated safety functions and performance standards to be identified. A verification scheme is then to be established to ensure that performance standards are – and continue to be – achieved. This means that the G1 safety argument for a sub-system can focus on whether the performance standards are met and safety functions delivered, without making a risk-based argument about each hazard. This supports clearer articulation of the overall platform safe to operate argument, while avoiding forcing the person with safety responsibility for the sub-system to try to make decisions about risk without knowing the full context. Local risks can still be addressed in the G2 argument, but in most cases will not form the most important part of the overall argument.

Def Stan 02-904 provides a definition of 'Safety Critical Items' that allows such items to be identified. Essentially, the scope of 'Safety Critical Items' comprises those items that provide mitigation against, or whose failure could cause or substantially contribute to, a loss of multiple lives associated with a Key Hazard Area, failure of life support systems for divers, or platform-level effects that could lead to severe damage or loss of the platform. Ideally, the safety functions performed by a Safety Critical Item and the associated performance standards would be known in advance of choosing a design. In the future, this may become more usual: the NATO ANEP-77 Naval Ship Code is expected to be updated to include a definition of 'essential safety function' which would align well to the concepts in Def Stan 02-904. Where this is not the case (e.g. when considering the safety case for operating or replacing a legacy system), it may be necessary to reverse-engineer the safety functions by examining the system design and the hazards involved, to identify those items that could lead to the type of losses described above, and to determine what role they might play in an accident sequence.

2.4. Complex Systems approach – IEC 63187

The Safety Critical Items approach described above is expected to be useful at the level of acquiring individual systems elements. It allows for some negotiation between the equipment level and the platform level, to agree what safety functions will be performed, and whether hazards will be mitigated at the platform or equipment level. However, when dealing with platforms built from complex systems, or themselves incorporated in higher-level capabilities, a more encompassing approach is needed.

The International Electrotechnical Commission (IEC) is currently developing a standard for the defence sector that addresses safety from a systems engineering viewpoint. *IEC 63187 – Systems engineering – System safety – Complex systems and defence programmes* will be based on ISO/IEC/IEEE 15288, the systems engineering life cycle standard (ISO, 2023). It takes the approach of augmenting the ISO/IEC/IEEE 15288 processes used to manage complex systems, to make them appropriate for safety-critical systems (Ricque et al. 2022, Inge et al., 2023). While many safety standards are designed primarily to be applied internally by an organisation (e.g. IEC 61508 (IEC, 2010), or between an acquisition organisation and a prime supplier (e.g. Def Stan 00-056 (MOD, 2023a)), IEC 63187 is being written for more complex supply chains, where multiple stakeholders will supply different system elements at different levels in the systems hierarchy.

 https://doi.org/10.24868/11189

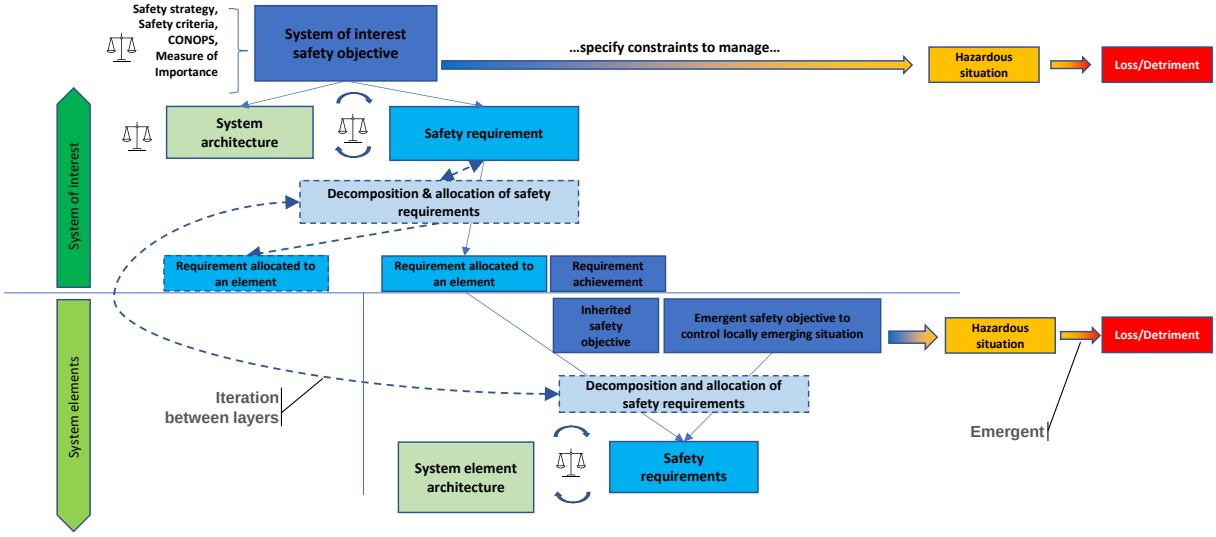

Figure 2. Cascade of safety objectives and requirements in IEC 63187 (Inge et al., 2023)

The new standard provides a framework for flowing safety objectives and corresponding requirements up and down the systems hierarchy (Figure 2). It has mechanisms to allow high-level platform safety objectives to flow down to system elements, and for objectives to be escalated to address interaction issues or hazards introduced by system elements that cannot be managed within the individual element. It also accounts for different elements being developed according to different life cycles and at different periods in time, e.g. legacy or commercial-off-the-shelf equipment being incorporated into a bespoke new capability.

Applying a framework like IEC 63187 at the architectural level on major programmes will help set clear requirements at the system element level, which will then feed into the safety functions and performance standards required for Safety Critical Items. In the future, this should make the Safety Critical Item approach described in section 2.3 easier to implement and more effective, by making it easier to see how individual sub-systems or equipment contribute to an overall platform or capability safety case, and what they have to do to be safe.

Conclusions

Haddon-Cave emphasised Cullen's view that 'safety cases were intended to be an aid to thinking about risk, not an end in themselves' (Haddon-Cave, 2009). Competent safety engineers are currently a limited resource across industry and the MOD. To make best use of this resource, and to have the best chance of improving safety, we must focus our safety case efforts on the most important parts of the safety argument, where there is the greatest opportunity to reduce risk. This paper has presented practical opportunities to reduce the effort spent on producing safety cases for low risk, low complexity items. It has also explained how structured arguments and modular safety cases can focus safety case production effort, and introduced work on two new standards, Def Stan 02-904 and IEC 63187, that will facilitate this for more complex systems.

References

ACWG, 2021. *GSN Community Standard Version 3*. York: Safety Critical Systems Club, Assurance Case Working Group No. SCSC–141C.

Adelard, 2024. Claims, Arguments and Evidence (CAE).

Haddon-Cave QC, C., 2009. *The Nimrod Review: An independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 Aircraft XV230 in Afghanistan in 2006*. London: The Stationary Office.

Hoare, C.A.R., 1981. The emperor's old clothes. *Commun. ACM*, 24 (2), 75–83.

IEC, 2010. *IEC 61508 – Functional safety of electrical/electronic/programmable electronic safety-related systems*. International Electrotechnical Commission.

IEC, 2023. *ISO/IEC/IEEE 15288 – System and Software Engineering – System life cycle processes*.

 https://doi.org/10.24868/11189

International Electrotechnical Commission.

Inge, J., Potiron, K., Williams, P., and Ricque, B., 2023. IEC 63187: Engineering Safety into Complex Defense Systems. *In*: *Safety in an Agile Environment: The International Systems Safety Conference 2023*. Portland OR, USA: International System Safety Society.

Inge, J.R., 2007. The safety case, its development and use in the United Kingdom. *In*: *Proc. 25th International System Safety Conference*. Baltimore MD, USA: International System Safety Society, 725–730.

Inge, J.R. and Costello, G.T., 2008. End-to-end reviews: A new approach to providing assurance that a complex organisation is effectively managing safety. *In*: *Proc. 26th International System Safety Conference*. Vancouver BC, Canada: International System Safety Society.

Leveson, N., 2011. The use of safety cases in certification and regulation. *J. syst. saf. (Online)*, 47 (6).

MOD, 2023a. *Def Stan 00-056 – Safety management requirements for defence systems – Part 1: Requirements*. Glasgow: Ministry of Defence, Defence Standard No. 00-056 Part 1 Issue 8.

MOD, 2023b. *JSP 815 – Defence Safety Management System*. London: Ministry of Defence, Joint Service Publication No. 815.

MOD, 2023c. *Def Stan 02-904 – Surface Ships Safety Critical Items*. Glasgow: Ministry of Defence, Defence Standard No. 02–904.

MOD, 2024. *DSA02-DMR – Defence maritime regulations for health, safety and environmental protection*. London: Ministry of Defence, Defence Regulation No. DSA02–DMR.

Ricque, B., Joguet, B., Brindejonc, V., Semeneri, N., and Potiron, K., 2022. IEC 63187 : intégrer la sûreté de fonctionnement au sein de l'ingénierie système. *In*: *Congrès Lambda Mu 23 " Innovations et maîtrise des risques pour un avenir durable " - 23e Congrès de Maîtrise des Risques et de Sûreté de Fonctionnement*. Paris Saclay, France: Institut pour la Maîtrise des Risques.

SDF, 2022. *Peer review of safety cases*. Nuclear Industry Safety Directors Forum.

 https://doi.org/10.24868/11189