Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)
Edited by Mário P. Brito, Terje Aven, Piero Baraldi, Marko Čepin and Enrico Zio
 Published by Research Publishing, Singapore.
doi: 10.3850/978-981-18-8071-1_P073-cd

The many faces of safety cases

Tor Stålhane

NTNU, Trondheim, Norway. E-mail: stalhane@ntnu.no

Thor Myklebust

SINTEF Digital, Trondheim, Norway. E-mail: thor.myklebust@sointef.no

This paper discusses the input documents and project decisions that are important when developing a safety case. The discussion is based on interviews with seventeen companies – all engaged in building safety cases for commercial products. The majority of the companies are Norwegian and Swedish. However, we have also interviewed companies from Denmark, UK, USA and Turkey. We discuss issues such as when in the project to start developing a safety case, what are the important inputs needed, and what are the roles of the required standards. Some issues will not be included – e.g. AI systems. The main reason for this is that none of the companies we interviewed developed AI systems.

We also discuss important issues such as the purpose of the safety case, safety case maintenance and the role of reuse when developing a safety case. We will also discuss the relationships between safety case and trust case and how a safety case can be used in communication and to build trust in a system.

Our further work will focus on two important areas – traceability between the system and safety case, which is important in order to keep the safety case up to date during system changes – and the possibility of expanding the "case" idea to bridge the communication gap between software developers and customers or users.

Keywords: *safety case, confidence, trust, maintenance*

1 Introduction

The work presented here is done as part of the TrustMe project. The TrustMe project started on August 1, 2020 and will last until June 2024. The project's main goal is to develop a safety case for autonomous buses and a safety case for the public. Safety cases are important when establishing confidence in the technology. The long-term goal for the TrustMe project is a regular operation with passengers without an operator on board the bus. Trials started in Norway in 2022, where the operator of the self-driving bus is moved to a remote-control room for surveillance and possible control handover if incidents cannot be handled safely or correctly by the self-driving bus.

This paper is a walk-through of the ideas for a safety case that surfaced during seventeen group interviews performed via Microsoft Teams. All the interviewees work in companies which perform analysis of safety related systems. The interview results are used in two ways:

- As a survey – X out of seventeen companies uses method Y.
- As a focus group – bringing up new ideas that should be discussed

The intent of the paper is to show how varied the concept and use of a safety cases is. This holds both for its purpose, its use and how it is developed and displayed. Based on the collected data, we will discuss the following issues:

- How and when to make a safety case
- How does the safety case fit into the development project plan
- How can the safety case be used to build confidence and trust and to facilitate project communication?

The Safety Related Application Condition (SRAC), a logical consequence of developing a safety case, is mentioned in a few of the interviews, but we will leave it out in order to focus on the safety case only.

2 The companies and risk experts involved

We have interviewed 35 risk experts from 17 different companies and organisations. 13 of the companies and organisations are from Norway and Sweden; the others are from Demark, Turkey, UK and USA. In figure 1 we have shown the number of domain experts involved. The term "other domains" represent: Hydrogen, Carbone Capture Storage, windmills, tunnels and storm surge barriers.

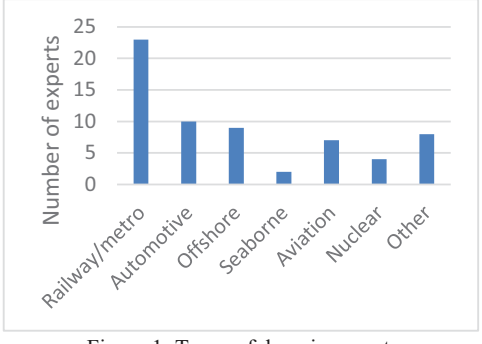

Figure 1: Types of domain experts

We have divided the experts into the following roles: assessor, safety case author, safety authorities and risk expert. About 50% percent of the experts have had more than one role. As seen from figure 2, more than 15 assessors and more than 15 safety case authors were interviewed.

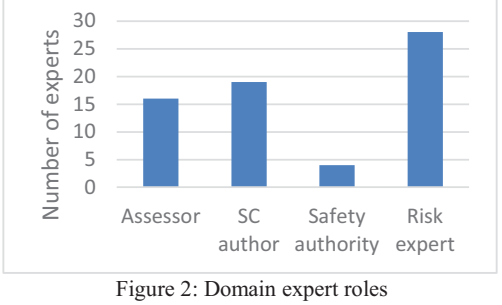

Figure 2: Domain expert roles

From the diagram in figure 4 we see that all the domain experts have a solid education with 28 of them having a master or PhD. It might be a weakness of the interviews that we have no person on our list that does practical work on a safety critical system – e.g., software developer or mechanical engineer.

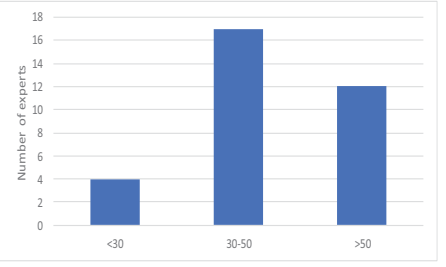

Figure 3: Age of domain experts

Most of the interview experts were above 30 years. Thus, most of them have more than five years of experience, although not necessarily related to safety cases.

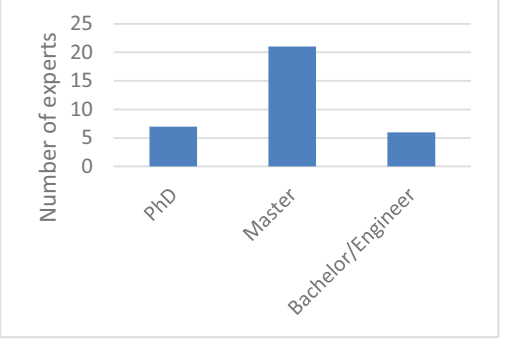

Figure 4: Domain experts' education

3 How and when to make a safety case

3.1 Interview results

In this section we will look at the following issues: what is the purpose of developing a safety case, and when do we start. Our advices are based on the input received from the interviewees.

The safety case is often a way to absolve the development company from any blame. In this case, the person who writes the safety case has the same function as the lawyer of the system provider. Ten of the seventeen companies interviewed – 56% – identified the cause for making a safety case. They all claimed that the safety case was needed to show that they had followed all relevant standards and regulations.

Most of the companies we interviewed – thirteen out of seventeen (76%) – wanted the development of a safety case to start as early as possible in the project, preferable at day one. An

alternative – mentioned by one company – was still to start as early as possible, but starting by developing a preliminary safety case. This preliminary safety case should then be updated in parallel with the project or process into a full safety case.

There should be a tight coupling between the safety case and the project plan. The reason is that the safety case developers will need information that must be provided by the development project. Thus, we need to start planning the safety case work together with the rest of the project planning process. Only five of the seventeen companies that we interviewed stated that writing a good safety case is time consuming.

Since different companies have different ways to do a safety analysis, they will identify different sets of safety-critical events and thus develop different safety cases – see for instance [Stålhane and Malm, 2014]. This is a challenge since different safety cases may lead to different conclusion – e.g., different levels of acceptance.

3.2 General advices

When writing the safety case, it is important to hit the right level of details – identify the critical issues and not focus too much on unimportant details. It is important to understand the relevant methods and standards used. Relevant safety case examples are important when writing a new safety case.

A lot of important decisions are made early in a development process. Thus, it is important to involve the assessor as early as possible and it is necessary to work on the safety case throughout the project. We should keep focus on the technology, technological questions, architecture and on quality management – within bounds of reason, quantum satis – but the main focus should be on safety. V&V (Verifications and Validation) results should be an integral part of the safety case but an excessive amount of extra V&V for the safety case could cause developer overload.

There are several ways to present a safety case – diagrams (GSN – Goal Structuring Notation), structured text, unstructured text, and spreadsheets – the latter mostly used for check lists. It is useful to look through old safety cases to see if there are any opportunities for reuse or at least some inspiration. However, there is danger if reuse of older safety cases gets too mechanistic. Only six of the seventeen companies we interviewed mentioned safety case reuse or using safety case patterns.

Writing a safety case is often seen as a lot of bureaucracy that takes a lot of resources and takes effort away from "real engineering work." In some cases, the safety case is used as a garbage can – a repository for everything that we may be concerned about, whether it is relevant or not, just to be on the safe side.

A final consideration is whether we should also include a negative safety case. A negative safety case should start with a claim that "The system is unsafe if…" What follows is the negative safety case. Even though this is not required by any standard, it will be a useful exercise both for future maintenance planning and for use of the system in new environments.

4 The safety case and the project

4.1 The safety case and the safety plan

In this section we will look at important input documents for a safety case – the client's safety requirements, the safety analysis and the development company's project plan, developed based on the customer's requirements, the safety analysis and the hazard log. Eight of the seventeen companies mentioned these documents. One of the companies stated that it was important to "prove" that they have done everything required by the applicable standards – proof of compliance. There is a tight coupling between the safety plan and the safety case. However, among the companies we interviewed, there were different opinions about the sequence of these two:

- Five out of seventeen companies stated that the safety plan must be developed first, based on the customer's safety requirements, since it shows how we will meet these requirements. The safety case will show that we have met each requirement.
- Eleven of the seventeen companies stated that the safety case must be developed first – partly based on the customer's safety requirements and partly based on the early safety analysis of the planned system and the hazard log. The claims, arguments and needed proofs of the safety case plus the relevant standards will be used to develop the safety plan.
- One company did not mention this issue

In all cases, we need to consider how the relevant standards will influence the safety plan and the safety case. The way we work – the company's development process – will also influence the safety case, especially the arguments and the evidence used. Thus, we also need to consider the relationships between the safety case and the company's methods for product development. Although not formalised in the standards, it is considered good practice to submit the safety plan to the Independent Safety Assessor / Risk Assessment Body (ISA/AsBo) before the final safety case is released. Whatever approach is used, it is important that the safety plan caters to future safety case needs.

4.2 Safety cases, DevOps and maintenance

Only eight of the seventeen companies we interviewed stated that they kept information on the relationships between product and safety case. Thus, in many cases, the safety case is not included in the change impact analysis. In addition, just as we maintain a software system, we need to maintain the related safety case(s) – safety cases of subsystems, products, items or equipment on which the system under consideration depends.

Most systems will change over time – new functionality, new hardware needed and new operator instructions. The hardware part of a safety case should also contain information related to necessary spare parts, although only one company mentioned this. Both hardware and related software – e.g., the operating system – may change over time.

We need traces between the safety case and system components. The traces will be needed in a change impact analysis throughout the project and later during maintenance. Both error corrections and system changes will require the ability to trace system changes to safety case components and the other way around. The trace info needs to be maintained. As a consequence of this, a safety case must be a living document. The increasing popularity of DevOps will increase the need for traceability since improving the customers' ability to give feedback will lead to more frequent changes.

For large systems, the final safety case will be built on top of many related safety cases – e.g., safety cases of any subsystems, products, items or equipment on which the system under consideration depends. Any change to a related safety case may lead to changes for the system's safety case. Often, the related safety cases are beyond the control of the personnel that have developed the main safety case. Thus, it is important to make sure that they are informed if those responsible for a subsystem do changes to their part of the system. An additional problem is that it is difficult to keep all the needed documents alive.

5 The many uses of the safety case

5.1 Building confidence in safety

The main purpose of a safety case is to build confidence in the claim that the product is safe for its intended purpose. Six out of seventeen companies stated that this was done by showing that they have followed all applicable rules and standards. The confidence required will to a large degree depend on relevant regulations and context. Not everybody agree that a safety case shows that the system is safe – one of the companies that we interviewed even claimed that a safety case proves nothing – it is just a description of how we worked. Five of the seventeen companies stressed the use of safety cases as a means for communication and building a safety culture – see section 5.3

The challenge is that we need to convince the safety case readers that we have done a good job – e.g., by documenting that the personnel that developed the system and the safety case had sufficient experience, knowledge, and training. In order to reach out to the safety case readers, it is also important that those who develop the safety case use a vocabulary well-known by the recipients.

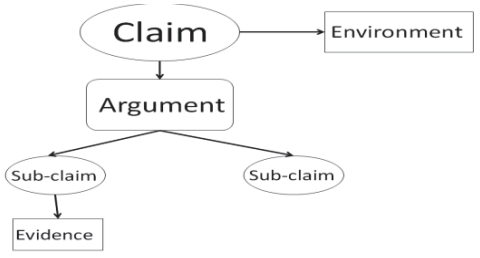

Figure 5: The main components of a safety case

Standards, together with claims, are important since we need to show that the claims comply

with the standards. The arbitrary structure of the safety case is a problem – there are several ways to decompose a claim into sub-claims – see figure 5. However, none of the interviewees had a solution to this problem. Whether a decomposition point of view is applied on a higher or on a lower level of the structure, seems to be a subjective decision of the users. In addition, we need to consider that there may be issues that cannot be analysed.

We should start with the claim – usually "The system is safe." We then go on to define the system and the system's operating environment. Then we need to write our claims. When these are agreed upon, we need to identify the relevant arguments, e.g., is "We have followed standard XYZ" a relevant argument. If it is not, then what is?

If relevant standards did not exist, it would be difficult for everyone to agree on all the requirements for the product. With the standards there are a minimum set of requirements – e.g., what all the sensors must achieve to claim compliance with the standards. It is much easier to require compliance with e.g., IEC 61508, than to write a large number of requirements. However, the safety standards are often difficult to use since the requirements are spread all over the document. It is thus difficult to know if all requirements are take care of in the safety case.

A safety case can be based on the system's requirements, the applicable standards and results of the safety analysis or a combination of these three. Safety case and safety standards are not independent. The process needs to comply with the standards, but different companies put different weights on these two components.

The RAMS (Reliability, Availability, Maintainability and Safety) department has to pull the necessary information from the project. Would it be better if those who already had the information available wrote the safety case? Note that the safety case preparation also may serve as an exploring trail by feeding detailed expectations to the design departments. Ownership of the safety case content by the technical disciplines is crucial.

5.2 Building trust in the product

The safety case vs. the trust case – what unite and what separate them? If we look at the trust case in the diagram in figure 6, we see that the part called "Learned trust" is close to a safety case. It focuses on how the system is designed and the users experience with the system. Part of the situational trust is also often included in a safety case.

Based on the trust model from [Hoff and Bashir, 2014], the TrustMe developed a trust case [Stålhane and Myklebust, 2022] – see figure 6. Trust has a situational (here and now) and a dynamic (learned) component. The situational component shows how the service provider will handle the "here and now" situations that may arise, e.g., emergencies and unexpected situations. The learned component shows how the service provider will collect and display information related to accidents and near-misses (track record), and shows how this info is used to increase trust. For a safety case, this side is usually not touched but the increasing popularity of DevOps may change this. The time component will also be involved if we use the "proven in use" argument.

In addition to the issues usually touched by a safety case, the trust model also includes assessment of task difficulty and how the system communicates with its environment – e.g., the users and operators. Both of these issues were mention by some of the companies participating in the safety case interviews and could thus be important.

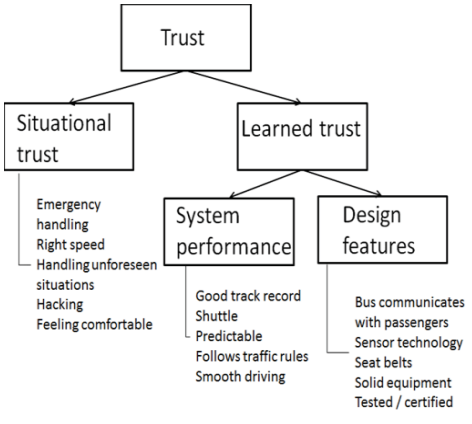

Figure 6: The trust model of the TrustMe project

Three of the companies we interviewed have already started mixing trust and safety by introducing the requirement that each argument and claim should be assigned a probability – how strongly do we believe this? The reliability / trust scores can then be combined using the method developed by Schaefer–Dempster [Schaefer–Dempster, 2018], giving the final trust or confidence in the safety case.

As part of a trust-survey, we asked the 55 participants to write down their own definitions of the difference between safety and trust [Stålhane and Myklebust, 2022]. The results are shown in table 1 – n indicates the number of "votes". Note that not all respondents answered this question. The important issue here is that while trust is based on feelings or belief, safety is based on an objective evaluation. Thus, we are of the opinion that a trust case and a safety case are mostly similar but while a safety case is based on evidence, a trust case is based on believes.

Table 1: Trust vs. safety

Trust	n	Safety	n
Feeling comfortable	15	Objective eval	14
Relying on somebody	14	Not in danger	13
Belief	9	Feeling safe	11
By choice	4	Handle accide	6
Subjective evaluation	4		
Building confidence	2		
SUM	48	SUM	44

5.3 Safety case as a means for communication

The safety case approach includes several benefits for improved communication both internally – inside the developing company – and externally – between the developing company and their customers and users. Thirteen of the seventeen (76%) companies we interviewed mentioned this. Relevant external stakeholders are customers, suppliers, assessors, certification bodies, notified bodies, assessment bodies (railway domain), authorities and the public [Myklebust and Stålhane, 2021]. Effective communication channels between all relevant stakeholders influence safety and in the end, the safety critical product or system.

Development of the safety case should be a continuous activity done in parallel with the development. However, only three of the companies we interviewed did this. The continuous activity will help developers and other project participants to have a common document and forum where they can discuss safety concerns and realizations problems throughout the project and also later during maintenance and upgrades.

One of the main reasons for using the safety case as an information tool is to show the developers why they need to do things in a certain way. E.g., include several safety aspects such as redundancy, safe state possibilities, handling reasonably foreseeable misuse and the ability of a functional unit to continue to perform a required function in the presence of faults or errors.

As part of the interviews, we have learned that it is especially important to make sure the readers' own vocabulary is used throughout the safety case. Through our development of safety cases since 2007 and SafeScrum [Hansen et al, 2018] since 2011, we have learned that communication between safety personnel and software developers is important for safety case development throughout the project. The authors of the safety case should have regular meetings with the developers and the RAMS team to ensure common understanding, to share knowledge and to establish the current status.

It is important to establish a culture for communication between the development project personnel and the management outside the project. Current safety standards are weak when it comes to requirements related to culture. Companies should check the safety culture regularly to ensure that the safety culture has an acceptable level and continuous to be on an acceptable level [Stålhane and Myklebust, 2022]. The safety case as a living document is also a good tool for the management to check the status of the project

6 Conclusions and further work

6.1 Conclusions

Our sequence of interviews with companies using safety cases so far has left us with the following main points – approximately half of the companies or more have indicated that:

- We need to start developing the safety case as early as possible – preferably as soon as the safety requirements are ready – 13/17
- A safety case is important for fostering efficient communication in the project – 13/17
- A safety case should be a living document. In order to do efficient maintenance of the

safety case we need a mechanism to link system and safety case claims and arguments – 8/17
- A safety case can be used to show that the project has adhered of all applicable standard and processes (10/17)
- The most important documents needed when writing a safety case are the system's definition, the hazard log and the safety plan – 8/17

In addition, some interesting points surfaced during the interviews. Mostly, they were just mentioned by one or two companies, but they are interesting because they touch upon problems we have seen in real life, both for safety and trust.

- It is important to use the glossary or terminology of the client, supported with customized argumentation models.
- Human factors are all too often not considered or only handled in a superficial way. We should e.g., consider the workload of the people operating the system
- We can mix trust and safety by introducing the requirement that each argument and claim should be assigned a probability – how strongly do we believe this?

6.2 Further work

We will try to find more companies to interview. Hopefully, this will confirm what we think is important and possibly add new issues to the safety case development process.

Only eight of the seventeen companies we interviewed stated that they kept information on the relationships between product and safety case. This is a serious problem since it implies that it is difficult to update the safety case when the system is updated. Thus, the safety case may no longer reflect the system's safety after a change. We need more research to find efficient ways to keep information links between system and safety case.

Several companies seem to use or at least acknowledge the safety case as a means of communication between safety experts and software developers. In our opinion, we could claim that the safety case bridge the knowledge gap between safety experts and software developers by clarifying why a certain action is needed. The important thing with a safety case is not "what" but "why". This could be taken a step further by construction communication cases which could be important for such issues as usability, installation and security. Improving communication will have two important benefits – the developers:

- Would understand why they had to do a certain thing.
- Could suggest a better solution, since they have a good overview of the existing code

Acknowledgment

This work has been supported by The Norwegian Research Council, Project name "TrustMe", project number: 309207 - IPOFFENTLIG19.

7 References

Hanssen, G.K., Stålhane, T. and Myklebust, T. (December, 2018) SafeScrum - Agile Development of Safety-Critical Software. ISBN 9783319993348. Springer.

Hoff, K.A and Bashir, M (September, 2014) Trust in Automation: Integrating Empirical Evidence on Factors That Influence Trust,

Myklebust, T., Stålhane, T. Gunnar D. Jenssen, G.D. and Wærø, I (2020) Autonomous cars, trust, and safety case for the public. RAMS Palm Springs USA.

Myklebust, T., Stålhane, T., Jenssen, G.D and Haug, I. (2021) TrustMe, we have a safety case for the public. ESREL, Angers France

Myklebust, T. and Stålhane, T. (December, 2021) Functional safety and proof of compliance. ISBN 978-3-030-86151-3, ISBN 978-3-030-86152-0 (eBook) https://doi.org/10.1007/978-3-030-86152-0. Springer International Publishing.

Myklebust, T. and T. Stålhane, T (2018): The Agile Safety Case. Springer

Stalhane, T. and Malm T. (2016) Risk Assessment – Experts vs. Lay People,

Stålhane, T. and Myklebust, T. (2021) Trust Case and the link to safety case. SAFE 9th International Conference on Safety and Security Engineering. Rome, Italy

Stålhane, T. and Myklebust, T (2022) Trust and acceptance of self-driving buses, ESREL, Dublin, Ireland

Pogodin, D. (2018) Features of Dempster–Shafer Theory Application in Planning of Construction Production, IOP Conf. Ser.: Mater. Sci. Eng. 463 042047