Goal-Based Safety Cases for Medical Devices: Opportunities and Challenges

Mark-Alexander Sujan[1], Floor Koornneef[2], and Udo Voges[3]

[1] Health Sciences Research Institute, University of Warwick, Coventry CV4 7AL, UK
m-a.sujan@warwick.ac.uk
[2] Delft University of Technology, TPM - Safety Science Group, P.O. Box 5015, 2600 GA Delft, The Netherlands
f.koornneef@tudelft.nl
[3] Forschungszentrum Karlsruhe GmbH, Institut für Angewandte Informatik, Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany
udo.voges@iai.fzk.de

Abstract. In virtually all safety-critical industries the operators of systems have to demonstrate a systematic and thorough consideration of safety. This is increasingly being done by demonstrating that certain goals have been achieved, rather than by simply following prescriptive standards. Such goal-based safety cases could be a valuable tool for reasoning about safety in healthcare organisations, such as hospitals. System-wide safety cases are very complex, and a reasonable approach is to break down the safety argument into sub-system safety cases. In this paper we outline the development of a goal-based top-level argument for demonstrating the safety of a particular class of medical devices (medical beds). We review relevant standards both from healthcare and from other industries, and illustrate how these can inform the development of an appropriate safety argument. Finally, we discuss opportunities and challenges for the development and use of goal-based safety cases in healthcare.

1 Introduction

In most safety-related industries the operators of systems have to demonstrate a systematic and thorough consideration of safety. In healthcare there is currently a differentiation between manufacturers of medical devices on the one hand and healthcare providers as users or consumers of such devices on the other hand. The current regulatory practice implies that the device manufacturers are responsible for determining acceptable levels of risk and for ensuring that the device is adequately safe for use in a specific context. However, the manufacturer usually has limited control over how devices are used in the operational context, and whether critical assumptions about aspects, such as training and maintenance are fulfilled. In addition, the healthcare service provider often has to integrate a number of different devices within their environment. The safety of the resulting system can only be assured if sufficient information from the manufacturer is provided.

F. Saglietti and N. Oster (Eds.): SAFECOMP 2007, LNCS 4680, pp. 14–27, 2007.

In this paper we show how the approach of a goal-based safety case can be used to analyse the safety of a medical device throughout its lifecycle, and to document the respective evidence used. Such an approach aims to overcome the limitations of current practice that results from the two disjoint regulatory contexts. It is also a first step towards investigating the feasibility of more complex system-wide safety cases.

Section 2 provides an argument for goal-based safety cases. In section 3, the regulatory context in the medical device area is described (using the UK as an example). The related standards and their dependence are presented in section 4. To explain the problem more closely, medical beds are used as an example for a medical device, and the outline of a safety case for this example is developed in section 5. Finally, section 6 concludes with a discussion of opportunities and challenges for the development and use of goal-based safety cases in healthcare.

2 Goal-Based Safety Cases

Argumentation is an important part of the development of safety critical systems. It provides information about why a system can be assumed to be sufficiently safe, and it may convey a measure of confidence. In many safety-critical industries such information is documented in a safety case. The purpose of a safety case can be defined as communicating a "clear, comprehensive and defensible argument that a system is acceptably safe to operate in a particular context" [1]. This definition reflects a goal-based approach, where the justification is constructed via a set of claims about the system, its behaviour, the process through which the system was produced, and about the safety case itself (i.e. the quality and the trustworthiness of the argument and the evidence produced) [2]. To support these claims specific evidence is produced. An essential component of goal-based safety cases is the argument that explains how evidence supports a particular claim. The argument makes explicit in the forms of rules and inferences the relationship between claims and evidence (see [3] for an extensive discussion of argument structure).

The use of goal-based arguments is now increasingly being reflected in standards, such as the UK Defence Standard 00-56 in its latest revision [4, 5]. Goal-based standards tell operators of systems what they need to achieve, rather than what kind of specific safety requirements they have to comply with. As technologies are evolving increasingly rapid, such an approach offers greater flexibility with respect to the use of novel beneficial technologies for which no corresponding assessment method has been defined in the standard, or practices that supersede outdated and costly development and assessment techniques [5].

Many standards also mandate incremental or evolutionary approaches to safety case development [6], such as the above mentioned UK Def Stan 00-56. Such an incremental approach can include multiple issues of, for example, Preliminary Safety Case, Interim Safety Case, and Operational Safety Case. At the Preliminary Safety Case stage the safety argument defines and describes the principal safety objectives, the general approach to arguing safety, and the types of evidence anticipated [1]. As

the design progresses and more detailed information becomes available, the arguments are subsequently extended and refined.

Narrative accounts of safety justifications often make it difficult for the reader or assessor to follow the logical argument that relates evidence to the claim it is intended to support. In addition, multiple cross-references make such documents generally hard to read and difficult to communicate to stakeholders of different backgrounds. Graphical argument notations, such as Goal Structuring Notation (GSN) [7] or ASCAD [8] explicitly represent the key elements of any safety argument, and their relationships. Tools have been developed that facilitate the construction of such graphical representations (SAM, ASCE) [9, 10]. With these tools the construction and the communication of safety cases is greatly facilitated [11].

3 The Regulatory Context

In many European healthcare systems there is a differentiation between manufacturers of medical devices on the one hand, and healthcare providers as users or consumers of such devices on the other hand. In general, manufacturers have to provide evidence that their devices are tolerably safe for a particular use in a specific environment. Healthcare providers, on the other hand, are audited to ensure that the care they provide meets national standards. A part of this is the requirement to utilise only previously certified medical devices. In this section we illustrate the certification process of medical devices and the audit of healthcare providers using the UK environment as an example [12].

The UK Medical Devices Regulations 2002 (MDR 2002) [13] implement a number of European directives relevant to the certification of medical devices. The definition of what constitutes a medical device is broad and comprises devices as diverse as radiation therapy machines, syringes and wheelchairs. The Medicines and Healthcare Products Regulatory Agency (MHRA) acts as the *Competent Authority* overseeing the certification of medical devices. *Notified Bodies* of experts provide evaluation of high and medium risk medical devices undergoing certification to the Competent Authority. The Medical Devices Directive [14] specifies essential requirements that have to be met by any device to be marketed in the EU. It provides classification rules based on the risk that medical devices pose, as well as conformity routes that specify different ways of manufacturer compliance with the essential requirements based on the class of the medical device under consideration. For most medical devices compliance with the requirements is demonstrated not through an explicit argument, but rather through either a self-certification process (lowest risk class) or through the compilation of specified evidence, including general product description, results of the risk analysis, and testing and inspection reports.

Apart from issuing instructions for use, the manufacturer of common medical devices has little influence on the way the devices are actually used in practice. More importantly, the manufacturer does not have detailed information about the specific environment and the processes within which the device will be operated within a

particular healthcare provider's setting. In complex systems this is a serious cause for concern, as in this way the possible interactions between system components and interactions with the environment as well as the system's particular history will not have been accounted for. It is reasonable, therefore, to expect healthcare providers to demonstrate that the services they are providing are acceptably safe. Such a demonstration should make use of data supplied by the manufacturers.

At present, healthcare providers are audited through a process that involves a number of diverse actors and agencies. The annual review of, for example, NHS Trusts is undertaken by the Healthcare Commission. The aim of this review is to establish whether Trusts comply with standards set out by the Department of Health [15]. These standards include aspects of safety, but are generally broader focussing also on issues such as financial health. During the annual review it is assessed whether Trusts have basic mechanisms in place, such as risk management and incident reporting, and whether there is evidence of continuous progress, e.g. learning from incidents. The data is collected throughout the year and includes national data about Trust activities, information from local organisations, documentation provided by the Trust, meetings with patient groups, as well as data from brief site visits conducted by assessors. The focus is on collecting indicators of (in the case of safety) safe practices, and accordingly the recommendations are targeted at specific issues, such as the fact that patients may not receive appropriate levels of nutrition, or that lessons learned from incidents are not shared among directorates.

In conclusion, therefore, within the UK regulatory context, both manufacturers of medical devices and healthcare service providers are regulated and are required to provide some kind of evidence that their devices and the services they provide are acceptably safe. However, in most cases there is no formal argument, and the two regulatory contexts (certification and audit) show little integration. This implies that assumptions and dependencies may not be documented properly, that interactions and unintended consequences of changes may go unnoticed, and that there are no formal notions of issues such as confidence in the evidence or diverse evidence to mitigate possible uncertainty (see e.g. [16] for an attempt of a corresponding formalism).

4 Relevant Standards

Safety of medical devices is regulated in about a thousand standards. In Europe, over two hundred of these are harmonised and provide a technical interpretation of the essential requirements of the MDD [14]. The main standard for electrical medical systems is the IEC 60601 series, which is now well underway in its 3rd edition revision process that started in 1995. The IEC 60601-1 series consists of Part 1: general requirements for basic safety and essential performance [17], and a number of collateral standards IEC 60601-1-XY on EMC, radiation, usability, alarm systems, environment, and physiologic closed-loop controllers. In addition, a particular standards series IEC 60601-2-YZ addresses specific requirements for particular systems, e.g. anaesthetic systems (60601-2-13), ECG monitoring equipment

(60601-2-27), ultrasonic equipment (60601-2-37) and screening thermographs (60601-2-56). Since 2005, the 3rd edition includes an update of all "tried and true" requirements of the 2nd edition and introduction of solutions now possible due to the availability of new technology. It also formalises the introduction of "Risk Management" by integration of standard ISO 14971 [18], see below, in order to make the standard less dependent on rapid growth in technology, and because there is more than "tried and true" requirements listed in the standard. Thus, it can be stated that medical electrical equipment shall remain single fault safe or the risk shall remain acceptable.

ISO 14971, entitled Application of risk management to medical devices, is now in its 2nd edition. This industry standard requires that the manufacturer shall establish, document and maintain throughout the life-cycle a process for identifying hazards associated with a medical device, estimating and evaluating the associated risks, controlling these risks, and monitoring the effectiveness of the controls. The manufacturer needs to demonstrate that the residual risk of using the medical system is acceptable. This is assumed when compliance to the requirements specified in IEC 60601-1 is demonstrated, but the bottom line is that the acceptability of a risk is determined by the manufacturer's policy for risk acceptability. Compliance to the performance of the risk management process is done by inspection of the risk management file (RMF). The RMF is the set of records produced by the risk management process, and remains the manufacturer's property that is not available to users.

Programmable Electrical Medical Systems (PEMS) are addressed in IEC 60601-1 clause 14 and regarding software elaborated in IEC 62304: Medical device software - Software life cycle processes [19]. Note that "it is recognized that the manufacturer might not be able to follow all the processes identified in clause 14 for each constituent component of the PEMS, such as off-the-shelf (OTS) software, subsystems of non-medical origin, and legacy devices. In this case, the manufacturer should take special account of the need for additional risk control measures."

Safety-critical human factors requirements are addressed in the standard IEC 62366 - Medical devices - Application of usability engineering to medical devices [20]. It states that "the manufacturer shall establish, document and maintain a usability engineering process to provide safety for the patient, operator and others related to usability of the human interface. The process shall address operator interactions with the medical device according to the accompanying documents, including, but not limited to transport, storage, installation, operation, maintenance, and disposal."

Other mainly technical standards exist for laboratory equipment, image processing systems and information networks, but these are not further elaborated in this paper. The standards of safety of medical systems have gone through a major revision process since 1995. This 3rd edition process will end with the implementation of the last collateral standard by about 2009.

The whole set of standards on safety of medical systems puts the manufacturer in the position of the decision maker of risk acceptance. The underlying assumption is

that a) the medical system will be used by "laymen", and b) the manufacturer defines normal use. All risk data about hazards, associated risks and acceptance criteria regarding a particular medical system has been elicited with adequate resources using inside information, and is recorded in the RMF. However, the professional user in a health care institution is left empty handed when they combine two medical systems in one configuration. [21, 22]. It is here that safety cases might help relevant parties to improve the understanding of risk assessment and control of operational risks related to the use of medical systems.

5 Development of the Top Level Argument for Medical Beds

5.1 Medical Beds

A medical bed is possibly the most stubborn medical device with respect to risk identification, control and management. Based on risk minimisation, an optimal bed may in many ways be bad for the patient as well as for the person providing care, because the bed will be higher than preferable for the patient, and lower than necessary for appropriate use of lifting and handling techniques. The patient is at risk due to gravitational forces (height) and opportunities for entanglement. Falling out of bed or getting strangled between bedrails are real harm scenarios. Occupational health and safety is relevant to nursing staff in particular: staff is at risk because of incorrect handling of the patient, wrong height of the bed, absence or wrong use of lifting aids, etc., potentially leading to serious back injury and disability to work. Hazards associated with medical beds include e.g. electricity, mechanical, electromagnetic interference and software failures in the bed motion control system. Mechanical hazards include e.g. entrapment, moving beds bumping into walls or other objects, bed instability, a collapsing component, and falls.

The medical bed (see Fig. 1) is defined in the particular standard IEC 60601-2-52 as a "device for which the intended use is sleeping or resting that contains a mattress

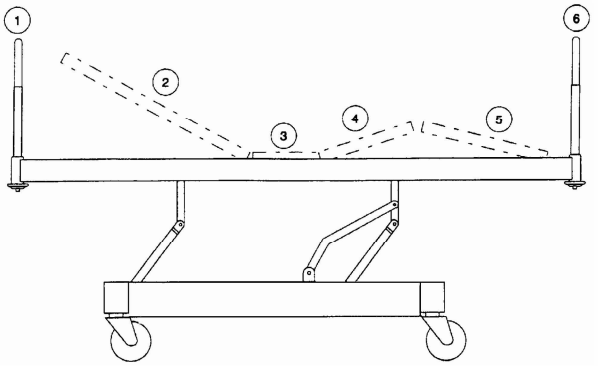

Fig. 1. Schema of a medical bed (from IEC 60601-2-52)

support platform. The device can assist in diagnosis, monitoring, prevention, treatment, alleviation of disease or compensation for an injury or handicap. A bed lift or a detachable mattress support platform in combination with a compatible non-medical bed as specified by the manufacturer is also considered a medical bed. Excluded are devices for which the intended use is examination or transportation under medical supervision (e.g. stretcher, examination table)." [23] Medical beds are meant for patients being defined as "person undergoing a medical, surgical or dental procedure, or disabled person". Maintenance of medical beds includes cleaning before reuse for another patient.

The standard identifies five distinct application environments, see IEC 60601-2-52:

1. Intensive/critical care provided in a hospital where 24 hours/day medical supervision and constant monitoring is required and provision of life support system/equipment used in medical procedures is essential to maintain or improve the vital functions of the patient.
2. Acute care provided in a hospital or other medical facility and medical electrical equipment used in medical procedures is often provided to help maintain or improve the condition of the patient.
3. Long term care in a medical area where medical supervision is required and monitoring is provided if necessary and medical electrical equipment used in medical procedures may be provided to help maintain or improve the condition of the patient.
4. Care provided in a domestic area and medical electrical equipment is used to alleviate or compensate for an injury, or disability or disease.
5. Outpatient (ambulatory) care which is provided in a hospital or other medical facility, under medical supervision and medical electrical equipment is provided for the need of persons with illness, injury or handicap for treatment, diagnosis or monitoring.

The use context of medical beds is important also because opportunities for managing operational risks differ.

5.2 General Top-Level Structure

A safety case essentially attempts to demonstrate that:

- The system under consideration is acceptably safe to enter service (in the UK this usually implies that safety risks are broadly acceptable or have been reduced as low as reasonably practicable).
- Arrangements are in place to ensure that the system will remain acceptably safe throughout its lifecycle.
- The structure and content of the safety case, and the process by which the safety case is produced and maintained are adequately trustworthy to produce a sound and convincing argument.

A common approach to demonstrate that the system under consideration is acceptably safe and continues to be so, is to argue that adequate safety requirements

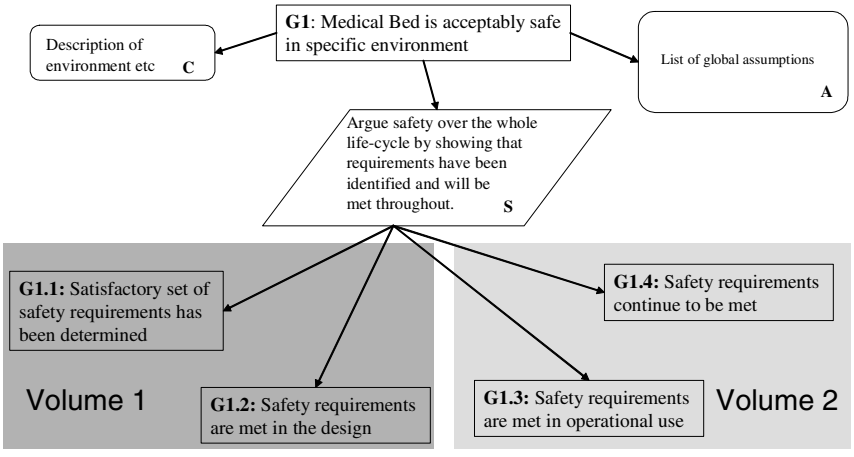

Fig. 2. Structure of the top-level argument and distribution of goals between manufacturer (Volume 1) and service provider (Volume 2)

have been established, that the safety requirements are met in the design, and that they continue to be met throughout all stages of the lifecycle of the system (see, for example, the objectives specified in the UK Defence Standard 00-56 [4]).

A well documented and frequently discussed example of a safety case formulated in GSN is the Eurocontrol RVSM Pre-Implementation Safety Case [24]. Here, the argument relies on four principle claims:

- Safety requirements have been determined and are complete and correct.
- The safety requirements are realised in the concept.
- The safety requirements are realised in the respective national implementations.
- The switch-over period (i.e. the transition from conventional vertical separation to the reduced vertical separation) is adequately safe.

One of the high-level safety requirements relates to continued safety throughout operational life:

- The performance and reliability of the system shall not deteriorate in service.

The arguments are then organized in such a way that for each it is demonstrated that sufficient direct evidence is available, and that this evidence is sufficiently trustworthy.

In principle, a similar general approach can be taken to demonstrate that medical devices, or more specifically medical beds, are adequately safe. We can argue that (see figure 2):

1. G1.1 A satisfactory set of safety requirements has been determined.
2. G1.2 Safety requirements are met in the actual design of the medical device.

3. G1.3 Safety requirements are met in operational use.
4. G1.4 Safety requirements continue to be met throughout the lifecycle of the medical device.

We can adopt the strategy taken in the RVSM Pre-Implementation Safety Case of arguing for each goal that there is sufficient direct evidence available, and that this evidence is sufficiently trustworthy.

Compared to current practice, where everything is addressed to and within the responsibility of the manufacturer, it is clear that healthcare service providers will have to provide some input. Objectives G1.3 and G1.4 exceed the control of the manufacturer. This crucially includes maintenance, as many serious accidents with technological systems relate to failures in the transition from maintenance mode to operational use and vice versa. Similarly, it cannot be assumed that the manufacturer can adequately manage operational risks through safety requirements that are met in the medical device. Rather, the service provider needs to demonstrate that arrangements are in place that satisfy assumptions made and ensure ongoing safe use and maintenance.

5.3 Outline of the Safety Case Structure

The four top-level goals G1.1 – G1.4 described above are then broken down until sufficient evidence has been provided that they are fulfilled, and an argument has been made that the evidence itself is sufficiently trustworthy.

G1.1 (A satisfactory set of safety requirements has been determined) is satisfied by arguing that a set of safety requirements has been identified, and that the safety requirements are complete, consistent and correct. The key strategy followed is the argument that relevant standards have been identified and addressed, and that the identified risks are sufficiently mitigated by the derived safety requirements. This is done by demonstrating that a risk management process according to ISO 14971 has been followed, and by providing the respective evidence.

G1.2 (Safety requirements are met in the design) is satisfied by arguing that the physical and functional properties of the medical device comply with the safety requirements, that procedure and training design comply with the safety requirements, and that any residual risks are tolerable.

G1.3 (Safety requirements are met in operational use) is satisfied by demonstrating that the guidance provided and assumptions made by the manufacturer of the medical device are taken into account by the service provider during operational use, that a hazard identification and risk assessment has been conducted by the service provider, and that risks have been sufficiently controlled through the specification of any required additional safety requirement.

G1.4 (Safety requirements continue to be met throughout the lifecycle) is satisfied by reference to the quality and safety management system of the service provider, and by demonstrating that adequate communication channels between service provider, device manufacturer and corresponding regulatory authorities have been established.

The Appendix provides a sketch of a previous preliminary argument development that was created during a session of the EWICS Medical Devices subgroup.

5.4 Arguing Safety over the Product Lifecycle

As discussed in section 4, the relevant device standards currently address the responsibilities of the device manufacturer. Standards at the organizational level of service provision, such as the UK Department of Health Standards for Better Health [15], are not concerned with medical devices apart from the requirement to use only those devices that have been certified. In the case of medical beds – a Class I medical device – a process of self-certification on part of the manufacturer is all that is required. This implies that decisions about levels of acceptable risks and detailed documentation of hazards considered remain with the manufacturer.

When healthcare providers assemble different devices to create a system within their environment, the safety of the resulting system needs to be assured. To this end the service provider needs to ensure that medical devices are installed according to the manufacturer's instructions for use, that appropriate maintenance is available, and that training and support to the operational staff is provided.

Apart from issuing instructions for use, the manufacturer has little influence on the way the devices are actually used in practice. The manufacturer does not have detailed information about the specific environment and the processes within which the device will be operated within a particular healthcare provider's setting. In complex systems this is a serious cause for concern, as in this way the possible interactions between system components and interactions with the environment as well as the system's particular history will not have been accounted for [12].

The structure that was chosen for the safety case in this paper attempts to bridge this gap by arguing the safety of the medical device throughout its lifecycle. Goals G1.1 and G1.2 (Safety requirements, and safety requirements met in the design) are clearly addressed to the manufacturer, while goals G1.3 and G1.4 (safety requirements are met in operation, and continue to be met) are addressed to the service provider. While in this respect – and quite reasonably – these two parts can be regarded as two volumes of the safety case, the requirements for each have now considerable impact on the structure and the content of the other.

In practical terms, the question arises how the two volumes could be sensibly separated into independent entities that can be produced by manufacturers and by service providers at different points in time. Here, the differentiation between different types of safety cases proposed in the CENELEC standard EN 50129 for railway applications (now to become IEC 62425) [25, 26, 27] could be a useful starting point. EN 50129 proposes three types of safety cases[1]:

- Generic Product Safety Case: provides an argument that a product can be used in safety-related applications; it specifies conditions that must be fulfilled in any safety-related application, where the product will be part of, and it provides descriptions of relevant characteristics of the product.

[1] Although EN 50129 is a standard for railway signaling systems, its definition of safety cases and their interrelationships are generic and are, in fact also applied outside the railway signaling field.

- Generic Application Safety Case: provides an argument that a configuration of products is adequately safe, without presupposing specific products.
- Specific Application Safety Case: provides an argument that a particular combination of specific components is adequately safe.

Each type of safety case specifies explicitly safety-related application conditions, the compliance with which has to be demonstrated in each safety case utilising that safety argument.

In the case of medical beds, the manufacturer would need to produce a *Generic Product Safety Case* (Volume 1 in fig. 2). As part of this, the device manufacturer needs to explicitly disclose all relevant assumptions made on the application, as well as all decisions regarding the acceptability of risks and the resulting mitigation. In addition, the manufacturer needs to demonstrate explicit forethought about the service provider's responsibility of ensuring safety during operation and throughout change. This entails, for example, documentation about appropriate procedures to operate the device, and the training needs of operators.

On the other hand, the service provider would need to produce a *Specific Application Safety Case* (Volume 2 in fig. 2), discharging the responsibility of demonstrating that the requirements established by the manufacturer as well as all assumptions explicitly made, are and continue to be satisfied during operation. In addition, as it is acknowledged that control of operational risks through safety requirements established by the manufacturer based on the device level is inappropriate, the service provider has to identify additional safety requirements based on their own operational environment. This is a very big change from the current practice of auditing that is carried out in order to collect indicators of safe practice.

Finally, to ensure continuing safe operation of the medical device in the operational environment, the service provider has to demonstrate that incidents are picked up, performance is monitored, the impact of changes is assessed, and crucially that effective communication channels to manufacturers and the relevant regulatory authorities are established. This responsibility is reflected on the part of the manufacturer by similar requirements that ensure that mechanisms for detecting and recording incidents and abnormalities are designed (where appropriate), and that arrangements are in place to receive and to react to data provided from the service providers or from regulatory authorities.

As a matter of speculation, we could envisage something similar to a *Generic Application Safety Case* being produced by professional bodies as guidance for the development of the specific safety cases to be produced by the service providers.

6 Opportunities and Challenges

6.1 Opportunities

The systematic consideration of safety through the development of goal-based safety cases has proven useful in industries such as aviation. The same benefits could be

expected in healthcare. Goal-based safety cases using graphical representation are easy to communicate, and can therefore address the variety of stakeholders of different (non-safety, non-engineering) backgrounds. Goal-based approaches also offer greater flexibility and are more suitable to incorporate novel technologies and methods.

In healthcare the disjoint regulation of device manufacturers and service providers has led to a situation where data from the two areas is usually not integrated, and where the device manufacturer defines both the normal operational context as well as acceptable levels of risk. Assurance of medical devices that have been put together by the service provider to form a particular system is hard to achieve. The development of a goal-based safety argument that demonstrates safety throughout the lifecycle of a device is an attempt at integrating data from manufacturers and service providers.

This approach could be a useful step towards whole system safety cases, e.g. for hospitals. This would be highly desirable as the individual device level usually is insufficient to assure safe operation, as the introduction of any device may have far-reaching unanticipated organizational reverberations.

6.2 Challenges

Healthcare does not have the same long tradition of reasoning about safety in systemic and explicit terms. Risk management in many healthcare organizations is still very preliminary and often includes only reactive approaches following incidents.

There is a split in the regulation between device manufacturers (certification) and service providers (audit). It is not clear who would be responsible for delivering such a safety case. Even if the safety case were split into separate volumes for manufacturers and service providers (e.g. along the lines of EN 50129 as proposed above), we may expect serious regulatory confusion as to which body is responsible for setting specific standards and requirements.

Many medical devices by themselves are not critical and do not require the development of a full safety case. However, in their specific application they may contribute to critical failures. The complexity of whole system safety cases needs to be addressed in future, as well as the process of integration of device manufacturers and service providers in the development of safety arguments.

Acknowledgments. Part of the presented work is based on discussions and work conducted within EWICS TC 7, Subgroup on Medical Devices. The safety argument was produced using the free academic license of the Adelard Safety Case Environment. We are grateful to Odd Nordland for input concerning EN 50129.

References

1. Kelly, T.: A Systematic Approach to Safety Case Management. In: Kelly, T. (ed.) Proc. of SAE 2004 World Congress (2004)
2. Bishop, P., Bloomfield, R., Guerra, S.: The Future of Goal-Based Assurance Cases. In: Proc. Workshop on Assurance Cases, pp. 390–395 (2004)
3. Toulmin, S.: The Uses of Argument. Cambridge University Press, Cambridge (1958)

4. DS 00-56 Issue 3: Safety Management Requirements for Defence Systems, Ministry of Defence (2004)
5. Kelly, T., McDermid, J., Weaver, R.: Goal-Based Safety Standards : Opportunities and Challenges. In: Proc. of the 23rd International System Safety Conference (2005)
6. Kelly, T., McDermid, J.: A Systematic Approach to Safety Case Maintenance. Reliability Engineering and System Safety 71, 271–284 (2001)
7. Kelly, T.: Arguing Safety, DPhil Thesis, University of York (1998)
8. Bloomfield, R., Bishop, P., Jones, C., Froome, P.: ASCAD – Adelard Safety Case Development Manual, Adelard (1998)
9. McDermid, J.: Support for safety cases and safety argument using SAM. Reliability Engineering and System Safety 43(2), 111–127 (1994)
10. Emmet, L., Cleland, G.: Graphical Notations, Narratives and Persuasion: a Pliant Approach to Hypertext Tool Design. In: Proc. of ACM Hypertext (2002)
11. Chinneck, P., Pumfrey, D., McDermid, J.: The HEAT/ACT Preliminary Safety Case: A case study in the use of Goal Structuring Notation. In: 9th Australian Workshop on Safety Related Programmable Systems (2004)
12. Sujan, M., Harrison, M., Pearson, P., Steven, A., Vernon, S.: Demonstration of Safety in Healthcare Organisations. In: Proc. Safecomp 2006, Springer, Heidelberg (2006)
13. Medical Devices Regulations 2002. The Stationery Office Limited, London (2002)
14. European Council: Council Directive 93/42/EEC of 14 June 1993 concerning medical devices. Official Journal L 169, 12/07/1993, pp. 0001 – 0043 (1993)
15. Standards for Better Health, UK Department of Health (2004)
16. Bloomfield, R., Littlewood, B.: On the use of diverse arguments to increase confidence in dependability claims. In: Besnard, D., Gacek, C., Jones, C.B. (eds.) Structure for Dependability: Computer-Based Systems from an Interdisciplinary Perspective, pp. 254–268. Springer, Heidelberg (2006)
17. IEC 60601-1 – Ed. 3.0 – Medical electrical equipment – Part 1: General requirements for basic safety and essential performance. IEC Geneva (2005)
18. ISO 14971:2007 – Application of risk management to medical devices. ISO Geneva (2007)
19. IEC 62304 – Ed. 1.0 – Medical device software – Software life cycle processes. IEC Geneva (2006)
20. IEC 62366 – Ed. 1.0 – Medical devices – Application of usability engineering to medical devices. Draft. IEC Geneva (2006)
21. 2nd EWICS MeD Workshop, Edinburgh (unpublished report) (2004)
22. Moore, S.: Integrating the Healthcare Enterprise - IHE NA 2007 Connectathon Fact Sheet (2006) Retrieved from (accessed 2007-03-19) www.ihe.net/Connectathon/upload/NA_2007_Connectathon_Fact_Sheet_1.pdf
23. IEC 60601-2-52 – Ed. 1.0 – Medical electrical equipment – Part 2-52: Particular requirements for basic safety and essential performance of medical beds. Draft. IEC Geneva (2006)
24. RVSM Pre-Implementation Safety Case, Eurocontrol (2001)
25. CENELEC EN 50129 – Railway Applications – Safety related electronic systems for signaling, CENELEC Brussels (2003)
26. Nordland, O.: Safety Case Categories – Which One When? In: Redmill, F., Anderson, T. (eds.) Current issues in security-critical systems, pp. 163–172. Springer, Heidelberg (2003)
27. Kelly, T.: Managing Complex Safety Cases. In: Proc. 11th Safety Critical Systems Symposium, Springer, Heidelberg (2003)

Appendix: High-level argument structure for demonstrating the safety of medical beds