Evidence:

Supplements to: Using safety cases in industry and healthcare

A pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare

December 2012

Identify Innovate Demonstrate Encourage

This research was commissioned and funded by the Health Foundation to help identify where and how improvements in healthcare quality can be made. The views expressed in this report do not necessarily represent the views of the Health Foundation.

This research was managed by:

Jonathan Riddell Bamber, Research and Development Manager
The Health Foundation

jonathan.bamber@health.org.uk
020 7257 8000

Project team

Prof Robin Bloomfield, Centre for Software Reliability, City University, London; Adelard LLP, London

Dr Nick Chozos, Adelard LLP, London

Dr David Embrey, Human Reliability Associates, Wigan

Jamie Henderson, Human Reliability Associates Wigan

Dr Tim Kelly, Department of Computer Science, University of York

Dr Floor Koornneef, Safety Science Group, TU Delft, Netherlands

Alberto Pasquini, Deep Blue Research & Consulting srl, Rome, Italy

Dr Simone Pozzi, Deep Blue Research & Consulting srl, Rome, Italy

Dr Mark-Alexander Sujan, Warwick Medical School, University of Warwick

Contributing authors

Dr George Cleland, Adelard LLP, London

Ibrahim Habli, Department of Computer Science, University of York

Dr John Medhurst, Human Reliability Associates, Wigan

Contact

Dr Mark-Alexander Sujan
Warwick Medical School
University of Warwick
Coventry, CV4 7AL

m-a.sujan@warwick.ac.uk

Evidence: Using safety cases in industry and healthcare is published by the Health Foundation, 90 Long Acre, London WC2E 9RA

Contents

Abbreviations

ALARP – As low as reasonably practicable

ASCE – Assurance and safety case environment

COMAH – Control of major accident hazards regulations

CAE – Claims, argument and evidence

CQC – Care Quality Commission

DSCN – Data set change notice

ETA – Event tree analysis

FDA – Food and Drug Administration (USA)

FMEA – Failure mode and effects analysis

FMECA – Failure mode, effects and criticality analysis

FTA – Fault tree analysis

GSN – Goal structuring notation

HAZOP – Hazard and operability study

HSE – Health and Safety Executive

IEC – International Electrotechnical Commission

ISO – International Organization for Standardization

MAPP – Major accident prevention policy

PCA – Provider compliance assessment

QRP – Quality risk profile

SCR – Safety case regulations

SCS – Safer Clinical Systems

SMS – Safety management system

Supplement A:

Safety case use in commercial aviation

Simone Pozzi and Alberto Pasquini;
Deep Blue Consulting & Research, Rome, Italy

This is one of a series of supplements to the report: *Using safety cases in industry and healthcare: a pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare*

To access the report and the other supplements, please visit www.health.org.uk/safetycasesreport

Contents

Supplement A:

Safety case use in commercial aviation

1 Introduction

The major challenge to be tackled in the civil aviation world is to maintain a stable safety performance, while coping with a relevant increase in traffic. That means having not more than the current number of incidents/accidents but on a higher number of flights. Such a need is strongly driven by the public perception of air transport safety: 'The air transportation industry's future viability may well be predicated on its ability to sustain the public's perceived safety while travelling'.[1]

The current safety level has been achieved thanks to a shift from reactive safety approaches to proactive ones. Reactive safety approaches were mainly reacting to undesirable events by prescribing measures to prevent their recurrence. Such approaches had, as a basic building block, the compliance with increasingly complex regulatory requirements, in order to ensure that minimum safety standards were met. Using these methods, the accident rate levelled off in the late 1970s, with the fatal accident rate in the vicinity of 10-6 (one fatal accident per million flights). Proactive approaches were adopted in the following two decades, by adding other factors to the compliance with regulatory requirements (which remained in place as a framework, on the basis of International Civil Aviation Organization (ICAO) standards and recommended practices – SARPs). These factors, included among others:

- safety culture and management's commitment
- implementation of Standard Operating Procedures (SOPs)
- checklists and briefings
- reporting systems and just culture
- competent investigations indentifying systemic deficiencies
- safety and human factor (HF) training for operational personnel
- sharing safety lessons among companies and states
- systematic safety oversight and performance monitoring.

One of the main distinguishing features of the aviation industry is certainly its transnational nature, whereby communities and organisations that are spatially distant come to interact with each other. This applies to the integration of national systems into a seamless network (like in the European Union (EU)), but also the multiplicity of actors that need to coordinate and interact for the air transport industry to deliver its services. Some of the main actors are:

- aviation professionals (flight crew, cabin crew, air traffic controllers, maintenance engineers)
- aircraft owners and operators
- military users of the airspace
- manufacturers
- aviation regulatory authorities (eg, Civil Aviation Administration (CAA), European Aviation Safety Agency (EASA), Agence pour la Sécurité de la Navigation Aérienne en Afrique et à Madagascar (ASECNA))

– industry trade associations
– professional unions and associations (eg, International Federation of Airline Pilots' Assocications (IFALPA), International Federation of Air Traffic Controllers' Associations (IFATCA))
– international aviation organisations (eg, ICAO).

An increasing integration and closer coordination among all these actors is at the core of innovation programmes like SESAR in the EU[2] or NextGen in the USA.[3] As described in the next section, such a trend has also modified the EU regulatory framework, by moving all the regulatory activities under the European Aviation Safety Agency (EASA) for both the airborne and the ground part (including Aviation Traffic Services (ATS) and Aerodrome operations).

As for the use of safety cases, the aviation community is divided, with two markedly different approaches for aeronautical products (ie the airborne part) and for Air Traffic Management services (the ground part). For aeronautical products, the approach is based on the granting of a type certificate for an approved design (or certificates of airworthiness to individual aircraft). It is in general a compliance-based approach centred on demonstrating satisfaction of detailed, prescriptive specifications. For Air Traffic Management (ATM) services, the approach has historically been performance-based. It is centred on providing evidence that a particular performance level will be obtained. For air navigation services, this evidence is often provided in the form of safety cases. For this reason, the remainder of this section will focus on the ATM domain, indicating where information is also applicable to the airborne part.

2 Regulatory context and best practice

The regulatory context in the civil aviation industry covers three main levels:

– **International Civil Aviation Organisation (ICAO)**: the international body setting standards and recommended practices for the safe and orderly development of international civil aviation.
– **European Aviation Safety Agency (EASA)**: founded in 2002, it ensures safety regulation for both the airborne part (airworthiness) and the ground part (including Aviation Traffic Services and Aerodrome operations). Airworthiness was previously supervised by the Joint Aviation Authority, consisting of 29 national aviation authorities from the European Civil Aviation Conference (ECAC) area. EUROCONTROL was instead providing the regulatory function for Aviation Traffic Services and Aerodrome operations.
– **National aviation authorities**: civil aviation administrations are the state body responsible for implementing the legislative and regulatory provisions for aviation safety.

The ICAO Annex 11[4] places an obligation on the ATM providers to ensure the safety of air traffic, in respect of those parts that fall under their managerial control. Implicit to such obligation is the 'burden of proof', by which service providers have to demonstrate positively that the relevant safety requirements and regulations are satisfied. The demonstration of compliance is not restricted to the current safety performance level, but it also encompasses planned changes: 'A certified provider of air traffic services shall notify the national supervisory authority of planned safety related changes to the provision of air traffic services'.[5]

The latest safety-related regulation was recently proposed by EASA, who published a draft regulation 'on safety oversight in air traffic management and air navigation services'.[6] The draft regulation builds on previously published regulations as detailed in the following quote.

> *'Pursuant to Regulation (EC) No 216/2008, the Commission, assisted by the European Aviation Safety Agency (hereinafter referred to as 'the Agency'), is required to adopt the relevant implementing rules to provide a set of safety regulatory requirements for the implementation of an effective air traffic management (ATM) safety oversight function. Article 8b of Regulation 216/2008 requires these implementing rules to be developed based on the*

regulations adopted under the Single European Sky I. This regulation is based on Regulation (EC) No 1315/2007 of 8 November 2007 on safety oversight in air traffic management and amending Regulation (EC) No 2096/20056.'

Previous regulations (listed in (Commission Regulation (EC), 2005) have instead identified and adopted the right provisions of the EUROCONTROL Safety Regulatory requirements (ESARRs).[7,8,9] The EUROCONTROL requirements are mandated by the same EU Regulation No. 2096/05, hence they are not repeated in the EASA draft regulation, but remain as part of the regulatory context.

The following table summarises the most relevant safety regulations and standards.

Regulation and requirement number	Extended title
EASA Draft Commission Regulation	On safety oversight in Air Traffic Management (ATM) and air navigation services
EU Regulation No. 691/2010	Laying down a performance scheme for air navigation services and network functions
EU Regulation No. 2096/05	Laying down common requirements for the provision of air navigation services
ESARR 1	Safety oversight in ATM
ESARR 3	Use of safety management systems by ATM service providers
ESARR 4	Risk assessment and mitigation in ATM

In the above context, safety management is defined as 'that function of air traffic services which ensures that all safety risks have been identified, assessed and satisfactorily mitigated, and that, secondly, a formal and systematic approach to safety management will maximise safety benefits in a visible and traceable way'.[10]

A typical safety management system (SMS) may be composed of the following functions:

- legislation, regulations, standards and practices
- incident reporting
- emergency debrief and documentation
- oversight: including safety survey, auditing, monitoring of performance
- safety assessment of changes to system/service.

Safety cases apply to the latter two categories, as they are meant to demonstrate the safety of:

- an ongoing service: falling under the safety oversight function
- a substantial change to that service: under the safety assessment function.

Whenever substantial changes are concerned, the safety case shall be produced when 'the severity assessment […] determines a severity class 1 or a severity class 2 for the potential effects of the hazards identified' or when 'the implementation of the changes requires the introduction of new aviation standards'.[11]

3 Development and drivers

Safety cases have acquired relevance in the civil aviation industry in the last decade, following the need of civil aviation administrations to implement a system for safety oversight and monitoring. Within Europe, EUROCONTROL has taken a leading role in supporting the use of safety cases for operational and regulatory changes. EUROCONTROL has developed guidance material[12] and example safety cases for ATM concepts.[13,14]

The purpose of the safety case is defined as follows.

'Broadly, the safety case is the documented assurance (i.e. argument and supporting evidence) of the achievement and maintenance of safety. It is primarily the means by which those who are accountable for service provision or projects assure themselves that those services or projects are delivering (or will deliver), and will continue to deliver, an acceptable level of safety.'[15]

The main driver for the safety case is the need to clearly summarise the processes and outcomes of safety activities to enable oversight by the competent authority (eg, the national CAA). Safety cases are used to collate all the safety-related evidences, to provide a structured index of safety activities and results. Oversimplifying for clarity's sake, such a need is engendered by the technical complexity of the systems under scrutiny, which would prevent any form of external independent oversight, if not organised according to a clear structure. The safety case should support an external reader in following 'the logical reasoning as to why a change (or on-going service) can be considered safe'.[15] In the same line the ICAO Safety Management Manual states: 'The documentation should be prepared in sufficient detail so that anyone reading it will be able to see not just what decisions were reached, but what the justification was for classifying risks as acceptable or tolerable'.[16]

Among the major obstacles to a widespread safety case adoption, the following aspects can be mentioned:

- Technical complexity of the safety case itself.
- Human resources to complete and review the safety case.

On the service provider's side, the safety responsible persons should master the technical aspects and complexities of the safety case itself, knowing how to structure it, the underlying methodology, the suitable notation and wording. A recurring problem is the maintenance of safety cases, which tend to be produced for major changes, but can be hardly updated in case of other subsequent changes to the same system components. Service providers may end up producing a new safety case every time one is required, either because the safety case scope has changed, or just because the safety case responsible is a different person.

On the CAA's side, the main obstacles are related to:

- the availability of human resources to deeply analyse and understand all the safety cases
- the required technical competences to understand the actual contents of the safety cases.

4 Types of safety cases and content

Two main types of safety cases exist:

Unit safety cases, which are used to demonstrate the safety of an ongoing service. A unit safety cases would include typically a safety assessment, to show that service/system is predicted to be safe, together with the results of safety audits, surveys and operational monitoring, to show that, up to that point in time, it has been safe. It should also demonstrate that processes are in place to ensure that all future changes to the system will be managed safely through, *inter alia*, project safety cases.

Project safety cases are used to demonstrate the safety of a substantial change to that service (and/or underlying system). They would normally consider only those risks created or modified by the change and rely on an assumption (or evidence from the corresponding Unit safety case) that the pre-change situation is at least tolerably safe. Project safety cases are used to update, and are usually subsumed into, unit safety cases.

A typical safety case contains the following elements.[17]

Aim	What the safety case is trying to show - this should be directly related to the claim that the subject of the safety case is acceptably safe
Purpose	Why is the safety case being written and for whom?
Scope	What is, and is not, covered?
System description	A description of the system/ change and its operational/ physical environment, sufficient only to explain what the safety case addresses and for the reader to understand the remainder of the safety case
Justification	For project safety cases, the justification for introducing the change (and therefore potentially for incurring some risk)

Argument	A reasoned and well-structured safety argument showing how the aim is satisfied
Evidence	Supporting safety evidence to substantiate the safety argument
Caveats	All assumptions, outstanding safety issues, and any limitations or restrictions on the operation of the system
Conclusions	A simple statement to the effect that the aim has been satisfied, subject to the stated caveats

Evidence can be derived from three main sources:

- **Service experience of previous usage**: data from previous operational use. Backing evidence should also be produced to show that the environment from which the data were obtained is sufficiently similar to that to which the re-used product will be subjected, that adequate performance assessment and fault recording processes were in place when the product was originally deployed, and that the analysis of the outputs of those processes was adequate and properly carried out.
- **Verification and validation**: analysis of any proof of requirements satisfaction that is obtained from the design or other representation of the product, including models, prototypes, software source code, etc. It includes, for example, simulation formal proof, hardware reliability prediction, inspection, and software static and dynamic code analysis. It also includes the final testing phase in an environment which is as close as possible to the operational environment.
- **Compliance with standards**: adherence to a particular standard can be used to demonstrate compliance with Safety Requirements, but it will depend on the nature of the standard itself.

The case-based approach has been recently extended to other areas than safety, with the introduction of the business case, environment case, human factors case, standards and regulation case, security case, technology case.[18] The purpose of the case-based approach is to group information into a clear structure in order to systematically identify and treat issues and benefits under each of the case perspectives.

Each case brings the benefit of a different specialist view, but such a multiplicity also comes with the risk of fragmentation, which can lead to incoherent/ inconsistent results. The European Operational Concept Validation Methodology provides guidance on how to strengthen the relationship between the cases – for instance, by ensuring coherency and integration of key outputs, or sharing expertise among the different teams, or sharing data in the case process, etc.[19]

5 Discussion

The main open issues concerning the use of safety cases in the aviation world concern the following:

- Integration of the compliance-based approach with the performance-based one, in order to reduce the divide between the airborne part and the ground part.
- Trade-off between clarity and brevity. Given the complex technical contents of the typical aviation safety case, it is often hard to achieve the right balance between clarity and brevity. Most safety cases require a deep understanding of the system's technical aspects, an understanding that cannot be developed through the safety case itself. As a result, the safety case may be understandable only by someone who already knows the system in detail. In other words, it may be unfeasible to achieve what ICAO recommends: to prepare the documentation 'in sufficient detail so that anyone reading it will be able to see not just what decisions were reached, but what the justification was for classifying risks as acceptable or tolerable'.[16]
- Tension between self-demonstration and regulatory approval. EUROCONTROL writes that by complying with its primary purpose (ie assuring the service providers of the acceptable level of safety), the safety case 'should also provide an adequate means of obtaining regulatory approval for the service or project concerned'.[15] For the just mentioned trade-off between clarity and brevity, the safety case may

not constitute the optimal communication and coordination means between the provider and the regulator. Other coordination processes may need to be put in place to ensure a through discussion of safety-relevant issues, thus bypassing the safety case function with less systematic and structured processes. This may result in the safety case being mostly developed to demonstrate to oneself that the system is safe (coherently with its definition), but leaving the burden of the provider–regulator relationship on other means.

- Defining the right scope. Risk assessment activities start by the definition of 'the scope, boundaries and interfaces part being considered'.[20] This determination is highly difficult, requiring complex decisions and trade-off between width and breadth. For instance, the scope may be too large, thus resulting in a too complex safety case, or it may be excessively narrowed down, to the point that the right perspective on the interactions among the different system components is lost.
- Use of severity assessment to classify changes. The safety case shall be produced when 'the severity assessment […] determines a severity class 1 or a severity class 2',[11] which correspond to *accident* or *serious incident*. The next severity class is *major incident*.[21] Severity classes are based on objective facts and are indeed good means to classify events in retrospect, but can be hard to apply to predict the potential effects of a hazard. For instance, the difference between classes 1, 2 and 3 is mostly good or bad luck, as most *major incidents* might potentially engender an *accident*. As a result, whenever considering a change, the ATM service provider may have no clear criterion to separate changes requiring a safety case and the oversight authority's approval, from those requiring only an internal assessment.

6 Lessons and recommendations for healthcare

Considering the healthcare situation, safety cases can be fruitfully applied to make systematic and well-structured safety activities. This will bring immediate benefit to those preparing the safety case.

A well-organised process may be the basis for effective coordination among service provider, technology provider and regulator. However, such coordination has proved hard to obtain, with safety cases sometimes satisfying formal requirements. In other words, healthcare professionals should be wary of the risk of the safety case becoming a goal in itself, rather than a means for transparent safety analysis.

Another key risk is related to the excessive technical content of safety cases. Safety cases dealing with complex safety-critical organisations may easily result in a document that cannot be understood by anyone else than the writer itself, full of jargon and domain knowledge.

The two issues just mentioned would suggest the adoption of safety cases only in those contexts where there is a good level of safety maturity, both on principles and methods. Otherwise, accompanying actions should be put in place at the same time, to:

- advance safety maturity
- develop competence in safety methods and techniques
- adopt safety cases.

This should target all the various organisational actors, including regulators and oversight authorities.

A last caveat concerns the use of severity classes, which proved hard in the aviation world and could be even trickier in a less standardised world like healthcare.[22]

The best aviation resource that could be transferred to ATM is the safety case manual delivered by EUROCONTROL, in particular for its examples and for the end checklist.

7 References

1 ICAO. (2006). *Doc 9858 Safety Management Manual (SMM)*. Chapter 1, p2.

2 SESAR. (2007). *Definition of the Future ATM Target Concept - D3*.

3 Joint Planning and Development Office (JPDO). (2008). *NEXTGEN Integrated workplan: a functional outline*.

4 ICAO. (2001). *Annex 11 - Air Traffic Services*.

5 Commission Regulation (EC). (2005). *No 2096/2005 - Laying down Common Requirements for the provision of air navigation services*.

6 Commission Regulation (EU). (2010). *Draft Commission Regulation (EU) on safety oversight in air traffic management and air navigation services*.

7 EUROCONTROL. (2000). *ESARR 3 - EUROCONTROL Safety Regulatory Requirement. Use of Safety Management Systems by ATM Service Providers*.

8 EUROCONTROL. (2001a). *ESARR 4 - EUROCONTROL Safety Regulatory Requirement. Risk assessment and Mitigation in ATM*.

9 EUROCONTROL. (2009). *ESARR 1 - EUROCONTROL Safety Oversight in ATM*.

10 Commission Regulation (EU). (2005). *No 2096/2005 - Laying down Common Requirements for the provision of air navigation services*.

11 Commission Regulation (EU) (2010) *Draft Commission Regulation (EU) on safety oversight in air traffic management and air navigation services*. p9.

12 EUROCONTROL. (2006). *Safety Case Development Manual*.

13 EUROCONTROL. (2001b). *The EUR RVSM Pre-Implementation Safety Case*.

14 EUROCONTROL. (2004). *The EUR RVSM Post-Implementation Safety Case*.

15 EUROCONTROL. (2006). p6.

16 ICAO. (2006). Chapter 13, p11.

17 EUROCONTROL. (2006). pp9–10.

18 EUROCONTROL. (2010). *European Concept Validation Methodology (E-OCVM). Version 3*.

19 EUROCONTROL. (2010). p53.

20 Commission Regulation (EU). (2005), p333/25.

21 Commission Regulation (EU). (2005), p333/26.

22 Pasquini, A, Pozzi, S, Save, L, and Sujan, M-A (2010). Requisites for Successful Incident Reporting in Resilient Organisations. In E. Hollnagel, D. Woods & J. Wreathall (Eds.), *Resilience Engineering in Practice: A Guidebook*. Aldershot, UK: Ashgate.

Supplement B:

Safety case use in the automotive industry

Ibrahim Habli
University of York

This is one of a series of supplements to the report: *Using safety cases in industry and healthcare: a pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare*

To access the report and the other supplements, please visit www.health.org.uk/safetycasesreport

Contents

Supplement B:

Safety case use in the automotive industry

1 Introduction

With the growing complexity of, and reliance on, safety-related systems in the automotive industry, the development of a safety case is increasingly being adopted as best practice for providing assurance for safety, particularly for electrical/electronic items. Interest in safety cases is being driven by the emergence of the draft international standard ISO 26262 (Road vehicles – Functional safety) in which the progressive development of a safety case is a core compliance requirement.[1]

2 Regulatory context and best practice

Technical regulations and standards provide the regulatory framework in the automotive industry. Whereas compliance with technical regulations is mandatory, often as a result of enacting a standard into law, adherence to a standard is voluntary, typically to demonstrate compliance with best practice. In Europe, United Nations Economic Commission for Europe (UN-ECE) regulations[2] and European Union (EU) directives[3] provide the main technical regulations for type approval. Due to the degree of overlap between the UN-ECE regulations and EU directives, they are considered to be equivalent for the purpose of system and component type approval. Compliance with technical regulations is demonstrated through third party approval, involving audits by an independent government-appointed authority. The United States operates a system of self-certification, administered by the National Highway Traffic Safety Administration (NHTSA) against the Federal Motor Vehicle Safety Standards (FMVSS).[4] Of the current 44 testable FMVSS (30 vehicle level, 14 equipment level), NHTSA selects a random sample of vehicles from the marketplace and independently tests them to verify that the Original Equipment Manufacture (OEM) certification is valid. In addition, there are seven non-testable FMVSS, which are verified by visual compliance.

3 Development and drivers

The automotive industry has until recently mainly adopted IEC 61508 (the generic international standard for electrical and electronic systems)[5] or design instructions such as the Motor Industry Software Reliability Associations (MISRA) Guidelines for Safety Analysis of Vehicle Based Programmable Systems[6] as an example of best practice for the development, operation and maintenance of embedded electrical and electronic systems. However, in 2004, two national initiatives in Germany and in France decided to merge and submit a proposal to the ISO for an automotive-specific safety standard. This was accepted and a new ISO working group ISO/SC22/TC3/WG16 (26262) was convened in 2005. The proposed standard is essentially an adaptation of IEC 61508 with the generation of a safety case as a key requirement. The impending introduction of ISO 26262 will offer the OEM and suppliers an agreed industry standard for managing risk for electronic vehicle systems. However, the concept of a safety case is not readily well known to those working in the automotive industry. The current draft of the ISO 26262 standard (DIS 2010-07-09) defines a safety case as an:

'argument that the safety requirements for an item are complete and satisfied by evidence compiled from work products of the safety activities during development'.[1]

This definition explicitly acknowledges the key role of the argument within the safety case, as compared to earlier proposed definitions such as the definition of a safety case as a 'compilation of work products of all safety lifecycle phases' (BL 9 2007-07-20).7 The current draft of the ISO 26262 standard (DIS 2010-07-09) requires a phased development of a safety case 'in accordance with the safety plan' where the safety case 'should progressively compile the work products that are generated during the safety lifecycle'.[1]

4 Types of safety cases and content

Current practices in the development of automotive safety cases are tightly linked to the evidence (or work products) generated from an ISO 26262 compliant process.[8–11] A typical automotive safety case includes high-level argument strategies for justifying the absence of unreasonable risks by appealing to the satisfaction of safety requirements during and after product development. These safety requirements include pre-defined certification requirements for showing that the vehicle has been homologated against regulations. They also include safety requirements derived from the hazard analysis and risk assessment process, which specify targets and means for preventing or mitigating the identified hazards. The argument within an automotive safety case is typically structured as a hierarchy of tiers, which consider the following:

- **Safety goals**: top-level safety requirements defined to address hazardous events identified by the hazard analysis and risk assessment process
- **Functional safety requirements**: derived from the safety goals and define implementation-independent safety measures or behaviour, eg functional redundancy and emergency operation
- **Technical safety requirements**: define how the functional safety requirements should be implemented in such a form which can be allocated to hardware and software elements.

The argument is not limited to the assurance of design safety but also addresses production and through-life safety through the consideration of field servicing and means for evaluating and responding to incidents and accidents. The justification of the organisational aspects of safety management is also a key element of the safety argument, requiring evidence to support that a company maintains a safety culture and manages staff competency. Finally, the type and the degree of rigour of the evidence supporting the argument depend on the risk classes associated with the automotive system. These risk classes take the form of Automotive Safety Integrity Levels (ASILs). ASILs (determined in terms of severity, probability of exposure and controllability) specify '*safety requirements for achieving an acceptable residual risk*' as well as '*requirements for validation and confirmation measures to ensure a sufficient and acceptable level of safety being achieved*'.[1]

5 Discussion

Although the concept of a safety case was considered in earlier automotive safety guidelines, the draft international standard ISO 26262 has significantly increased interest within the automotive safety industry in how safety arguments and evidence should be generated, documented, reviewed and maintained for automotive systems. Currently, there is not a consensus on the real value of an automotive safety case, particularly when a safety process is compliant with ISO 26262. On the one hand, some are treating the safety case as a repository of the work products generated from the safety lifecycle phases. On the other hand, others are emphasising the role of the argument in showing how and why the work products (ie evidence) support the overarching claim that residual risks are acceptable. MISRA has recently launched an initiative to create guidance on automotive safety cases. It aims to provide supporting information on the structure and content of the argument, as well as worked examples of arguments and sample means for generating evidence.

6 Lessons and recommendations for healthcare

Evidence, in the form of healthcare-specific case studies and pilot applications, is important to convince the community of the value of a safety case. In particular, these case studies and pilot applications should show why and how a safety case adds value to the overall safety process. Understandably, many will initially regard the development of a safety case as a documentation exercise needed merely for compliance. However, clear and practical healthcare-specific guidance, supported by example safety arguments and evidence, should help in paving the way for a smooth introduction of the safety case concept in a way that ensures a consistent understanding of this concept.

7 References

1 International Organization for Standardization (ISO), ISO 26262: Road vehicles – Functional safety, DIS 2010-07-09, 2010.

2 World Forum for Harmonisation of Vehicle Regulations (WP 29), http://www.unece.org/trans/main/wp29/wp29wgs/wp29gen/wp29age.html (accessed on 4 March 2011).

3 Council of the European Communities, Council Directive 70/156/EEC of 6 February 1970 on the Approximation of the Laws of the Member States Relating to the Type-Approval of Motor Vehicles and their Trailers, http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31970L0156:en:NOT (accessed on 4 March 2011).

4 National Highway Traffic Safety Administration, Federal Motor Vehicle Safety Standards and Regulations, http://www.nhtsa.dot.gov/cars/rules/import/FMVSS/#SN101 (accessed on 4 March 2011).

5 International Electrotechnical Commission (IEC), IEC 61508: Functional Safety of Electrical / Electronic / Programmable Electronic Safety-Related Systems, 2010-04-30, 2010.

6 The Motor Industry Software Reliability Association (MISRA), Guidelines for Safety Analysis of Vehicle Based Programmable Systems, November 2007.

7 International Organization for Standardization (ISO), ISO 26262: Road vehicles – Functional safety, BL 9 2007-07-20, 2007.

8 Palin, R, Habli, I. *Assurance of Automotive Safety: A Safety Case Approach*, 29th International Conference on Computer Safety, Reliability and Security (SAFECOMP), Vienna, Austria, September 2010.

9 Habli, I, Ibarra, I, Rivett, R, Kelly, T. *Model-Based Assurance for Justifying Automotive Functional Safety*, 2010 Society of Automotive Engineers (SAE) World Congress, Detroit, USA, April 2010.

10 Dittel, T, Aryus, H-J. *How to 'Survive' a Safety Case According to ISO 26262*, 29th International Conference on Computer Safety, Reliability and Security (SAFECOMP), Vienna, Austria, September 2010.

11 Törner, F. *On Hazard Identification and Safety Cases In the Automotive Domain*, PhD Thesis, Chalmers University of Technology, Sweden, 2008.

12 Chen, D, Johansson, R, Lönn, H, Papadopoulos, Y, Sandberg, A, Törner, F and Törngren, M. *Modelling Support for Design of Safety-Critical Automotive Embedded Systems*, 27th International Conference on Computer Safety, Reliability and Security (SAFECOMP), Newcastle upon Tyne, UK, September 2008.

Supplement C:

Safety case use in the defence industry

Tim Kelly
University of York

This is one of a series of supplements to the report: *Using safety cases in industry and healthcare: a pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare*

To access the report and the other supplements, please visit www.health.org.uk/safetycasesreport

Contents

Supplement C:

Safety case use in the defence industry

1 Introduction

Safety case development and acceptance is an established practice for the assurance of safety in the procurement and operation of UK defence land, sea and air systems. For example, military aircraft, armoured fighting vehicles, submarines, and radar systems all have corresponding safety cases that justify their acceptable safety for operation in defined operational contexts. One notable distinction of the nature of safety cases in the defence domain is that safety is typically defined in terms of freedom from unacceptable risk of *unintentional* harm (given that there can be *intentional* harm).

Historically, many defence safety cases were limited to 'peacetime' operations (with the understanding that operation in war is inherently 'high risk'). However, this scope has become increasingly challenged in recent years, with greater attention to unacceptable and unintentional risk even during wartime operations. Also, the historical focus has been solely on unintentional harm inflicted by the Ministry of Defence's *own* equipment and personnel. However, there is growing interest in safety in terms of protection against known risks, *including* those posed by third parties.

Defence safety cases are not restricted to simply arguing the safety of equipment. They are also used to cover wider aspects of operation, incorporating safety arguments and evidence concerning all aspects of the defence lines of development (known by the acronym TEPIDOIL – training, equipment, people, infrastructure, doctrine, organization, information and logistics).

2 Regulatory context and best practice

The development and acceptance of safety cases as a means of assuring and certifying the acceptable safety of UK defence-related equipment has become widespread since the mid-1990s. There was a notable change in 1996 from Issue 1 to Issue 2 of Defence Standard (DefStan) 00-56 (Safety Management Requirements for Defence Related Systems)[1] to include the concept of safety cases as an essential part of the procurement of UK defence-related systems. In the time since Issue 2 of DefStan 00-56, the Ministry of Defence (MoD)'s own internal safety management standards (documented through Joint Service Publications - JSPs) have all gradually been updated to include the requirements for the development and acceptance of safety cases. Examples include JSP 430[2] requiring safety cases for ships and ship systems, JSP 553[3] requiring safety cases for military aircraft, and JSP 454[4] requiring safety cases for land systems.

The following quote from JSP 430 Issue 1 illustrates the MoD's internal requirements for the production of safety cases:

> *'Safety Cases are required for all new ships and equipment as a means of formally documenting the adequate control of Risk and demonstrating that levels of risk achieved are As Low As Reasonably Practicable (ALARP).'*

In addition, the following quote from DefStan 00-56 (now at Issue 4[5]) illustrates the requirement that the MoD places on its suppliers for the production of safety cases:

> *'9.1 The Contractor shall produce a Safety Case for the system on behalf of the Duty Holder. The Safety Case shall consist of a structured argument, supported by a body of evidence, that provides a compelling, comprehensible and valid case that a system is safe for a given application in a given environment.'*

3 Development and drivers

Safety case development and acceptance has been practised in some quarters of the MoD for many years, notably for the Nuclear Steam Raising Plants that power UK nuclear submarines (having adopted the approach from the civil nuclear domain). However, the original suggestion for the *widespread* introduction of safety cases in UK Defence came from Management Services (Organisation) Division Study No. 773 – Equipment Safety Assurance, dated March 1994.[6] This study recommended the implementation of a safety case regime in the procurement process, with the objective *'to provide a well organised and resourced justification of the acceptability of the System's safety'*. The study summarised that the purpose of the safety case was to identify the hazards and risks, indicate the safeguards, and describe the safety management system arrangements. It also recommended the modification of DefStan 00-56 to include safety cases (prompting the update to Issue 2 as described in the previous section).

The Secretary of State for Defence has a duty of care for the safety of MoD personnel under the provisions of the UK Health and Safety at Work Act. This duty of care is delegated through formal Letters of Delegation to key stakeholders in the MoD organisation. The MoD has chosen the safety case approach as one of its means of assuring that this duty is being discharged in its operations.

There has been recent criticism of the MoD's safety case regime (particularly focused on the air domain) in the Nimrod Review[7] conducted by Charles Haddon-Cave QC. The following concerns were raised (among others):

> *'It [the safety case regime] has led to a culture of "paper safety" at the expense of **real** safety.'*

> *'The current shortcomings of Safety Cases in the military environment include: bureaucratic length; their obscure language; a failure to see the wood for the trees; archaeological documentary exercises; routine outsourcing to Industry; lack of vital operator input; disproportionality; ignoring of age issues; compliance-only exercises; audits of process only; and prior assumptions of safety and "shelf-ware".'*

The Nimrod Review recommended that the practice of safety case development be retained by the MoD, but also that, 'Safety Cases should be renamed *"Risk Cases"* and conform in the future to the following six principles: Succinct; Home-grown; Accessible; Proportionate; Easy to understand; and Document-lite.' The recommendation to rename safety cases has not been adopted by the MoD. Internal reviews are still underway as to how to best respond to the Nimrod Review's other recommendations.

4 Types of safety cases and content

As described above, Defence Standard 00-56 is used by the MoD when procuring equipment to require suppliers to produce safety cases. By its own internal standards (JSPs), the MoD is required to have safety cases in place for all equipment. It is increasingly recognised that there is a distinction between the safety case that can be provided by an equipment supplier, and the wider safety case that can, and should, be in place for systems **in** operation. This has led to a distinction between equipment safety cases and operational (sometimes called operating) safety cases. For example, a military platform such as a Eurofighter Typhoon would have a military aircraft safety case developed as part of the equipment procurement and development process. However, in addition, there would also be a corresponding operational safety case for the Typhoon that takes the wider perspective of the platform in operation, its surrounding processes, support systems, logistical support, etc. This distinction between types of safety cases has been drawn out in the following recent safety guidance produced for the air domain:

> *'… the overall System Safety Case is likely to compose of the following three main components: a. Design Safety Case – managed by a Design Organisation under contract b. Project Safety Case – managed by the [MoD Procurement] Project Team; and c. User Safety Case – managed by the Operating Authority (User).'*[8]

Concerning the structure of a safety case, Defence Standard 00-56 contains an explicit requirement for structured arguments:

> *'9.5 The Safety Case shall contain a structured argument demonstrating that the evidence contained therein is sufficient to show that the system is safe. The argument shall be commensurate with the potential risk posed by the system and the complexity of the system.'*

Safety cases are documented and communicated between stakeholders by means of Safety case reports. Extra guidance on the acceptable form of a safety case report is provided in Part 2 of Defence Standard 00-56. For example, the standard recommends the following sections for any safety case report:

- executive summary
- summary of system description
- assumptions
- progress against the programme
- meeting safety requirements
- emergency/contingency arrangements
- operational information
- independent Safety Auditor (ISA) report
- conclusions and recommendations
- references.

In addition, the following guidance from Part 2 describes the recommended form of safety arguments communicated with safety case reports:

> *'9.5 All Safety Cases contain an implicit safety argument but this Standard requires an explicit argument; this is usually expressed in terms of a defined hierarchy of safety claims and sub-claims that are supported by a body of evidence.'*

5 Discussion and lessons for healthcare

There is a substantial body of experience with safety case development within the UK defence domain. The domain has now lived with the cross-service requirement for safety cases for more than fifteen years. As the Nimrod Review Report discusses, in this period there have been examples of poor safety cases[9] alongside positive examples. There is therefore much to learn from safety case practice in the UK defence domain. Positive features to observe include clear regulatory standards (such as Defence Standard 00-56), clear definitions, and the practice of producing explicit and structured safety arguments. Negative experiences to note include where safety case development appears to have become a 'paper exercise' through lack of

stakeholder engagement, where safety cases have not been used or allowed to influence and change practice, and where there has been insufficient distinct organisational support for the MoD's many roles with respect to safety cases (customer, operator, regulator, owner).

The defence domain is inherently high-risk. It could be argued that this is a shared characteristic with the health domain. In both domains we have to distinguish between preventable and unpreventable risk.

The defence domain, through the Secretary of State's declared responsibilities under the Health and Safety at Work Act, has been able to establish a clear regulatory requirement for safety cases that, for example, is now 'flowed down' to every equipment supplier. Questions as to the source and authority of the possible requirement for safety cases within the healthcare domain should be addressed.

The defence domain has evolved a clear understanding of the difference between equipment safety cases and wider operational safety cases, and that a hierarchy of safety cases is sometimes required in order to establish an overall safety case. It is important to consider the parallels to these concepts in the healthcare domain (eg, to consider how the overall safety case for an NHS trust might be decomposed in to supporting sub-cases covering different aspects of the trust's operations).

6 References

1 MoD, Defence Standard 00-56: Safety Management Requirements for Defence Systems, Ministry of Defence, Defence Standard, Issue 2, December 1996.

2 MoD, Joint Service Publication JSP 430: MoD Ship Safety Management, Ministry of Defence, Issue 1, January 1996.

3 MoD, Joint Service Publication JSP 553: Military Airworthiness Regulations, Ministry of Defence, 1st Edition, July 2003.

4 MoD, Joint Service Publication JSP 454: Procedures for Land Systems Equipment Safety Assurance, Ministry of Defence, Issue 3, July 2002.

5 MoD, Defence Standard 00-56: Safety Management Requirements for Defence Systems, Ministry of Defence, Defence Standard, Issue 4, June 2007.

6 MoD, Management Services (Organisation) Division Study No. 773: Equipment Safety Assurance, Ministry of Defence, March 1994.

7 Haddon-Cave, C. *The Nimrod Review: An independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 aircraft XV230 in Afghanistan in 2006*, The Stationery Office, ISBN 9780102962659.

8 MoD, *Safety Handbook Initial Release – A Guide to Good Safety Management Practice,* Ministry of Defence Director of Air Support, April 2010.

9 Kelly, T. *Are Safety Cases Working?*, Safety Critical Systems Club Newsletter, Vol. 17, No. 2, January 2008.

Supplement D:

Safety case use in the nuclear industry

Robin Bloomfield, Nick Chozos, George Cleland
Adelard LLP, London

This is one of a series of supplements to the report: *Using safety cases in industry and healthcare: a pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare*

To access the report and the other supplements, please visit www.health.org.uk/safetycasesreport

Contents

Supplement D:

Safety case use in the nuclear industry

1 Introduction

In this report, we present a review of the development and application of safety cases in the nuclear sector. The report is structured as follows.

Section 2 presents an overview of the legislative and regulatory context within which nuclear safety is assured. We discuss the stakeholders, the key legislation and guidance, and the role of the safety case within all these.

Section 3 presents the evolution of the nuclear industry in the UK, focusing on key events that have shaped safety regulation and the nuclear safety case.

Section 4 takes a more in-depth look at safety cases, considering approaches to their development and content. We take a particular focus on control and instrumentation (C&I) systems containing software.

Section 5 focuses on healthcare, considering what lessons there are for a potential medical safety case and for patient safety, based on the experiences of the nuclear domain.

Section 6 provides a concluding discussion.

2 Regulatory context and best practice

This section presents an overview of the UK and international regulatory context of nuclear safety, while focusing on the nuclear safety case; in particular, we identify the key legislation and requirements for nuclear installations, as well as the main actors and stakeholders, their roles and interactions that are based around the safety case.

The safety of nuclear installations in the UK is assured and managed by a system of regulatory control based on a *licensing process*. In this licensing process, a corporate body is granted a licence to use a site for specified activities (licensing requirements are discussed in detail by the Health and Safety Executive[7]) for a specified timeframe under specific conditions. As we shall see in this report, the 'nuclear safety case' plays an integral role in the acquisition and attainment of the licence through the planned life of the installation as it is the basis for both safety assurance and communication to the regulator that grants the licence.

The Health and Safety Executive (HSE) licensing guidance[7] explains that the 'licensee' is expected to:

> *'…make and implement adequate arrangements for the production and assessment of safety cases consisting of documentation to justify safety during the design, construction, manufacture, commissioning, operation and decommissioning phases of the installation'.*

The 'nuclear safety case' may relate to a nuclear site, a nuclear power plant, part of a plant, a plant modification, or a set of significant issues. Its overall purpose is to establish and present, in a written format, a case demonstrating that all requirements of the applicable legislation are met and that the system is, and will continue to be, acceptably safe to use within its environment throughout its lifetime.

2.1 UK nuclear safety regulation

Under UK law (the Health and Safety at Work etc Act 1974[5]) employers are responsible for ensuring the safety of their workers and the public. This responsibility is reinforced for nuclear installations by the Nuclear Installations Act 1965 (NIA).[4] Under the relevant statutory provisions of the NIA, a site cannot have nuclear plant on it unless the user has been granted a site licence by the Health and Safety Executive (HSE). The NIA stipulates that only a corporate body – a legally united body that can act as one individual, such as a registered company or a public body – can hold such a licence. This licensing function is administered on HSE's behalf by its Nuclear Directorate (ND). The legal regime is complemented by the Ionising Radiations Regulations 1999 (IRRs), which provide for protection of workers in all industries from ionising radiations, and by the generality of health and safety regulation which the ND also enforces on nuclear sites.

Her Majesty's Nuclear Installations Inspectorate (NII) (see Section 2.1.2.2) is the part of the HSE with responsibility for regulating the safety of nuclear installations in Great Britain. In April 2007, the UK nuclear security regulator, the Office for Civil Nuclear Security (OCNS) and the UK Safeguards Office (UKSO) joined the HSE. Together with the NII, they form the ND. NII inspectors use a set of Safety Assessment Principles (SAPs)[1] to guide regulatory decision making in the nuclear licensing process. Underpinning such decisions is the legal requirement on nuclear site licensees to reduce risks to *as low as reasonably practicable* (ALARP) (see Section 2.2.3 for more on ALARP), and the use of these SAPs should be seen in that context.

The term 'safety case' is used by the NII[1] to encompass the totality of a licensee's documentation to demonstrate high standards of nuclear safety and radioactive waste management, and any subset of this documentation that is submitted to the NII.

Figure D1: The UK regulatory framework for nuclear safety below presents an overview of the nuclear safety regulatory framework and identifies the key stakeholders and their roles.

Figure D1: The UK regulatory framework for nuclear safety

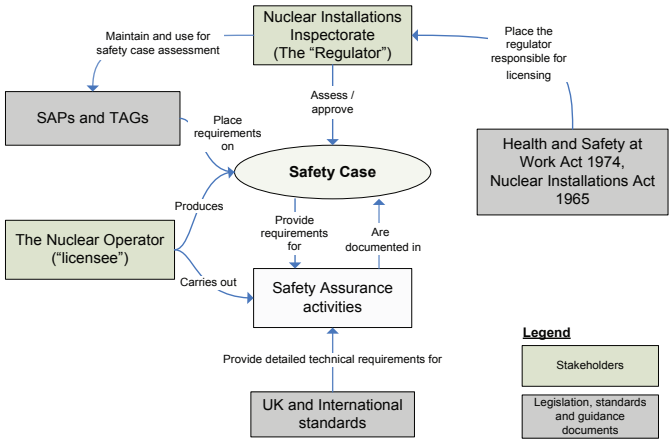

Within this framework, the regulator also carries out physical inspections at the plants to ensure that the conditions attached to the licence and all applicable regulations are complied with as described in the safety case.

The UK law that specifies the licensing requirements, the regulator, other stakeholders and their roles, the safety case, and its role in the licensing process, are discussed in more detail in the following sections.

2.1.1 Licensing requirements

2.1.1.1 Health and Safety at Work Act 1974

The Health and Safety at Work etc Act 1974[5] – also referred to as HASAW or HSW Act – is the primary piece of legislation that addresses occupational health and safety in the UK. The operators of nuclear facilities in the UK, like their counterparts in other industries and places of work in general, are required to comply with the requirements of the HSW Act.

The HSW Act imposes a duty on all employers to ensure, so far as is reasonably practicable, that persons not in their employment as well as those that are, are not exposed to risks to their health or safety as a result of the activities undertaken.

The HSE's approach to enforcement[18] of the requirements of the HSW Act is informed by the principles of:

– **proportionality** in applying the law and securing compliance
– **consistency** of approach
– **targeting** of enforcement action
– **transparency** about how it operates and what those regulated may expect
– **accountability** for its actions.

The HSE has a number of ways to help operators secure compliance with health and safety legislation, enabling it to take a proportionate approach in each case. The primary means are inspections and investigations, which are undertaken to gather information and analyse it to draw conclusions about the level of compliance with the HSW Act. Inspectors may offer licensees and duty-holders guidance and advice, both face to face and in writing. This could include providing a warning that, in the opinion of the HSE inspector, it is failing to comply with the law. Where appropriate, inspectors may serve improvement and prohibition notices, withdraw approvals, take action under the conditions attached to certain types of licences, or prosecute.

2.1.1.2 Nuclear Installations Act 1965

The Nuclear Installations Act was first established in 1959, following a public inquiry into the Windscale fire in 1957, the largest civil nuclear accident in the UK (discussed in more detail in Section 3.1).

The Act was amended in 1965 and is since then referred to as the Nuclear Installations Act 1965 (NIA). The NIA places particular requirements on the nuclear industry. Overall, it serves three purposes.[7]

– It requires that sites which are to be used for the operation of nuclear reactors are licensed.
– It provides for the control of the processes and the application of security measures associated with the enrichment of uranium and the extraction of plutonium or uranium from irradiated matter.
– It sets up a special legal regime to govern the liability of the licensees towards third parties for certain kinds of damage (primarily nuclear damage) caused by nuclear materials on, or coming from, their sites

The HSE is responsible for the first of those purposes, the licensing and inspection of sites. The other two are the responsibility of the appropriate Secretary of State, for sites in England and Wales, and Scottish Ministers for Scotland.

In the next section we take a closer look at the part of the HSE that is responsible for nuclear safety – the Nuclear Directorate and its Nuclear Installations Inspectorate.

2.1.2 The regulator

2.1.2.1 Health and Safety Executive – Nuclear Directorate

The nuclear licensing function, as mandated by the HSW Act and NIA, is administered on HSE's behalf by its Nuclear Directorate (ND).

The ND sets out conditions attached to the site licence that spell out general safety requirements to deal with the risks that exist on a nuclear site. Licensees will comply with these in different ways where arrangements and procedures will be in place to meet a licence condition. In most cases, a *safety case* is expected to demonstrate that the licensing requirements are met.

The ND seeks to stay current with and improve safety standards for work with ionising radiations at licensed nuclear installations. It does so through its licensing powers by reviewing and assessing safety cases and inspecting sites for licence compliance. It sets national regulatory standards and helps to develop international nuclear safety standards.

Inspections are carried out by the Nuclear Installations Inspectorate (NII), the enforcement section of the ND. The NII is presented in the next section.

2.1.2.2 Nuclear Installations Inspectorate

The Nuclear Installations Inspectorate (NII) has the overall responsibility for granting licences (and attaching conditions as appropriate). NII's inspectors are appointed under the HSW Act and they administer the NIA. They deal with nuclear and radiological safety issues at licensed nuclear sites.

The NII also monitors and regulates occupational, non-nuclear health and safety aspects, but this is done in conjunction with inspectors drawn from other parts of the HSE. The inspectors' activities involve technical assessment of the safety of

existing and proposed nuclear facility designs and their operational regimes, inspection of the implementation of the licensee's licence condition compliance arrangements and investigation of incidents and complaints.

The Safety Assessment Principles (SAPs) and the supporting Technical Assessment Guidelines (TAGs), which are used as a basis for these assessments, are described in more detail in the following sections.

2.1.3 The licensee

As we have previously mentioned, the nuclear operator is solely responsible for the development and maintenance of the safety case. For the operator, the safety case is not only a means of acquiring a licence, but, as we shall see later on, a driver for establishing a high level of safety.

Over the years, the UK nuclear operator has developed a safety culture which drives not only the engineering of safety and safe operation of the plant, but also a continuous effort in finding better ways to demonstrate safety. As such, with the combination of its safety culture and the regulatory context discussed here, the operator is willing to invest significantly in the production of a sound and compelling safety case by supporting research, contracting specialists and working closely with the regulator.

2.1.4 Requirements and guidance

2.1.4.1 Safety Assessment Principles

The HSE Safety Assessment Principles (SAPs)[1] are the primary principles that define the overall approach to be followed for nuclear installations in the UK. The SAPs have been revised in the past few years and have been brought into line with international guidance (from the International Atomic Energy Agency, which is discussed in Section 2.1.5). The SAPs stand as a framework for inspectors to make consistent regulatory judgements on nuclear safety cases. The principles are supported by Technical Assessment Guides (TAGs) (see Section 2.1.4.2), and other guidance, to further assist decision making by the nuclear safety regulatory process.

The contents of the SAPs are summarised in Table D1 below:

Table D1: Contents of the HSE Safety Assessment Principles

Principle	Description
Fundamental principles	These principles are founded in UK health and safety law and international good practice, and underpin all those activities that contribute to sustained high standards of nuclear safety.
Leadership and management for safety	Principles that form the foundation for the leadership and management for safety in the nuclear environment.
The regulatory assessment of safety cases	Principles applicable to the assessment of the production and nature of safety cases.
The regulatory assessment of siting	Principles applied in the assessment of a site, since the nature of a site can have a bearing on accident consequences.
Engineering principles	The major part of the SAPs and covers many aspects of the design and operation of nuclear facilities.
Radiation protection	Focus on the relevant principles of the Ionising Radiations Regulations 1999.
Fault analysis	Engineering principles concerning the detection and diagnosis of malfunctions in systems.

Principle	Description
Numerical targets and legal limits	Probabilistic targets to assist in making judgements regarding the tolerability of risk (Section 2.2.2) and the ALARP principle, which is discussed in Section 2.2.3.
Accident management and emergency preparedness	Principles on the procedures around dealing with incidents and accidents.
Other	Other principles regarding radioactive waste management, decommissioning, control and remediation of radioactively contaminated land.

The SAPs are very important as they provide not only principles for assessment for the regulator, but also set out the requirements for the duty-holder to adhere to. The contents of a nuclear safety case are most likely to be structured around the SAPs. It is also important to note that the principles explain what has to be done, but do not become prescriptive as to how these requirements will be met. They provide flexibility to the duty-holder to achieve what is expected by the SAPs. It is then a matter of the safety case to demonstrate adherence to the SAPs.

We focus on some of the contents of the SAPs later on in this report when discussing particular elements of a safety case in Section 4.

2.1.4.2 Technical Assessment Guides

The Technical Assessment Guides (TAGs) give additional guidance to the NII inspectors and assessors that go beyond the contents of the SAPs. In particular, the TAGs go into some level of technical detail, playing an advisory role, often used as reference for particular engineering activities. There are currently 79 TAGs covering areas such as hazard analysis, training and competence, human reliability analysis, radiological protection and others. There is also a TAG on the 'Purpose, Scope and Content of Safety Cases' which has been used to provide some of the content of this report.[6]

It is to the benefit of the operator to refer to these both for carrying out the engineering work and for producing the safety cases which will eventually be assessed on the basis of the SAPs and the TAGs.

2.1.4.3 UK and international standards

While the SAPs and the TAGs constitute the requirements and guidance for nuclear installations and their supporting safety cases in the UK, there is a set of international standards which must be complied with. Compliance requirements with these standards depend on the system functionality and technology.

Table D2 below identifies some of the key standards that influence the engineering and safety of systems in nuclear installations. Due to the abundance of nuclear safety-related standards, we have selected to present the ones that relate to control and instrumentation (I&C) systems – the systems that are used to monitor and regulate parameters such as temperature and pressure and perform control functions (such as emergency plant shutdown).

Table 3: Key international standards applicable to nuclear safety

International standard	Description
IEC 61508[33] – Functional safety of electrical/electronic/ programmable electronic safety-related systems	Central to the standard are the concepts of risk and safety function. The risk is a function of frequency (or likelihood) of the hazardous event and the event consequence severity.
IAEA NS-R-1[34] – Safety of nuclear power plants: design: safety requirements, September 2000	This publication establishes safety requirements that define the elements necessary to ensure nuclear safety. These requirements are applicable to safety functions and the associated structures, systems and components, as well as to procedures important to safety in nuclear power plants.
IEC 61513[32]– Nuclear power plants - Instrumentation and control for systems important to safety - General requirements for systems	Provides requirements and recommendations for the total I&C system architecture which may contain either of the following technologies used in I&C systems important to safety: conventional hardwired equipment, computer-based equipment or a combination of both types of equipment.
IEC 61226[24] – Nuclear power plants - Instrumentation and control important to safety - Classification of instrumentation and control functions	This standard establishes a method of classification of the information and command functions for nuclear power plants, and the instrumentation and control systems and equipment that provide those functions, into categories that designate the importance to safety of the function. The resulting classification then determines relevant design criteria.
IEC 60880[35] – Nuclear power plants - Instrumentation and control systems important to safety - Software aspects for computer-based systems performing category A functions	Provides requirements for the software of computer-based instrumentation and control (I&C) systems of nuclear power plants performing functions of systems important to safety (category A).

It is common that a safety case will attempt to demonstrate compliance with one or more standards. In order to do so, evidence must be provided that individual clauses of the standard have been met, or an acceptable justification in cases where a clause has not been met. The standards compliance approach to safety cases is discussed in more detail in Section 4.3.

Apart from these key publications, there is a plethora of standards that may be referred to for a variety of activities and systems that exist within a nuclear installation. For instance, ISO 9001, the international standard for quality management systems, is expected to be complied with both by the operator and contractors who are commissioned to carry out work for the operator. In addition, there are international agencies that aim to promote safety and security and streamline the approaches to safety across the planet by creating international forums, codes, practices and standards. These are discussed in the following section.

2.1.5 Nuclear Safety International agencies and regulation

The International Atomic Energy Agency (IAEA) is the international centre of cooperation in the nuclear field.[36] It was set up as the world's 'Atoms for Peace' organisation in 1957 as part of the United Nations. The IAEA works with its member states and other partners worldwide to promote safe,

secure and peaceful nuclear technologies. The IAEA's safety standards are not legally binding on member states but may be adopted by them, at their own discretion, for use in national regulations in respect of their own activities.

IAEA requirements are explicit in requiring a regulatory body to keep its principles, regulations and guidance under review from time to time, taking account of internationally endorsed standards and recommendations. The UK's goal-setting legal framework for health and safety does not apply IAEA requirements in a prescriptive manner, but they are reflected within the principles.

There are several advisory bodies that oversee the development of safety standards: the Advisory Commission for Safety Standards (ACSS); the Nuclear Safety Standards Advisory Committee (NUSSAC); the Radiation Safety Standards Advisory Committee (RASSAC); the Transport Safety Standards Advisory Committee (TRANSSAC); and the Waste Safety Standards Advisory Committee (WASSAC). Member states are widely represented on these committees.

NII is a member of the Western European Nuclear Regulators' Association (WENRA), which is dedicated to ensuring that all European Union countries and candidate countries with civil nuclear power stations, as well as Switzerland, have harmonised high levels of nuclear safety. To this end, WENRA has developed reference levels that represent good practices for existing civil nuclear power plants and is developing reference levels for radioactive waste management and decommissioning. Harmonisation requires there to be no substantial differences from the safety point of view in generic, formally issued, national safety goals, and in their resulting implementation on nuclear power station licensed sites. In the UK, the reference levels will be secured using a combination of: national laws; health and safety regulations; conditions attached to nuclear site licences; and the SAPs, TAGs and other forms of guidance used when granting nuclear site licences and in regulating licensees' activities.[2]

2.2 Purpose and objectives of the nuclear safety case

In the previous sections, we have presented the international and UK regulatory context and the key stakeholders that are there to ensure the safety of nuclear installations. In this section we focus on the safety case, and primarily the role it plays within this regulatory context.

The nuclear safety case is defined by the HSE[6] as:

> *'…the totality of documented information and arguments which substantiates the safety of the plant, activity, operation or modification in question. It provides a written demonstration that relevant standards have been met and that risks have been reduced as low as reasonably practicable (ALARP).'*

The licensee is legally responsible for the safety case. However, it is those employees of the licensee who have direct responsibility for delivering safety who should have 'ownership' of it. By this we mean an understanding of the safety case and limits and conditions derived from it.

2.2.1 Purpose of the safety case

According to the SAPs,[1] the purpose of the safety case is to establish and demonstrate in written form that the plant processes, activities, modifications, etc being proposed:

- are soundly assessed and meet required safety principles
- conform to good nuclear engineering practice and to appropriate criteria, standards and codes of practice
- are adequately safe during both normal operation and fault conditions
- are, and will remain, fit for purpose
- give rise to a level of nuclear risk to both public and workers which is ALARP. (The ALARP principle is discussed in more detail in Section 2.2.2 and Section 2.2.3)
- have a defined and acceptable operating envelope, with defined limits and conditions, and the means to keep within it.

The safety case also forms the basis for delivering safe operation. The analysis it provides of normal operation and possible accidents should identify the measures that need to be implemented to realise the required safety standards. These measures include: operating rules and instructions; examination, maintenance and testing requirements; minimum staffing levels in key areas (eg. control rooms); staff training needs; and emergency procedures.

In the following section we focus on the ALARP principle: a principle that is intrinsic to the definition and evaluation of safety in the UK.

2.2.2 Definition of risk

The HSE report 'Tolerability of Risk from Nuclear Power Plants'[39] defines risk in the context of operation of nuclear stations. In this report, the definition of nuclear risk is specified as: 'the probability that a specified undesirable event will occur in a specified period or as the result of a specific situation'.

Any risk calculation has to take into account both severity and probability of the event. In the context of risks relating to the operation of a nuclear power station, the risks of greatest interest are those associated with radiation – both to individuals and to society.

- **The risk to individuals**: this concerns the risk to the health of any particular individual, worker or member of the public. In this case, the harm may be either in the form of 'early effects' or 'late effects'. Early effects will occur if the radiation dosage is very high and will result in direct death. With regards to late effects, the greatest concern is cancer.
- **The risk to society**: societal risks are those that are wider than those to individuals, having consequences for the environment and infrastructure, such as loss of electricity, and the economy. For instance, in the case of the Chernobyl disaster, apart from the individuals that died or were affected by the fall-out, there were significant effects on the environment and consequently to the food chain, not only locally but also internationally.

These risks are considered both in normal operation of nuclear installations and in accident conditions and have to be reduced to a level of 'tolerability'. Tolerability does not mean 'acceptability' – it refers to a willingness to live with a risk to yield certain benefits so long as there is confidence that it is properly controlled. For this, design of nuclear installations and their supporting and safety systems is targeted towards minimising and controlling risks. Cost and rigour of activities must be proportionate to those risks. Calculations of risk, taking into account severity and likelihood, have to then demonstrate that the risk is appropriately mitigated to a tolerable level.

The next section discusses the 'ALARP principle': the principle that certain nuclear risks have to be demonstrated to be 'As Low As Reasonably Practicable'.

2.2.3 The ALARP principle

The interpretation of ALARP has been set into the context of 'tolerability of risk' as a clarification arising from the Sizewell B Public Inquiry (discussed further in Section 3.1). This section is based on a report for the UK Nuclear Safety Advisory Committee[22] (the HSE provides detail on the definition and interpretation of ALARP[19]).

The ALARP principle is based on the assumption that it is possible to compare *marginal improvements in safety* (marginal risk decreases) with the *marginal costs of the increases in reliability*. Nuclear risks may offer this possibility when they are quantified (ie in terms of event probability and of radiation releases), and when the failure rate improvements of the systems controlling the relevant events can be evaluated. This comparison of marginal variations does not in principle require a common measure, but simply that both risk and the marginal cost or efforts to improve reliability can be realistically assessed. This assessment can however be problematic, especially when design faults have to be taken into account.

This found its expression in the well-known 'carrot diagram' (see Figure D2), which has become the standard means for the exposition of the principle.

Figure D2: The ALARP principle: levels of risk are divided into three bands. Width of wedge represents level of risk.

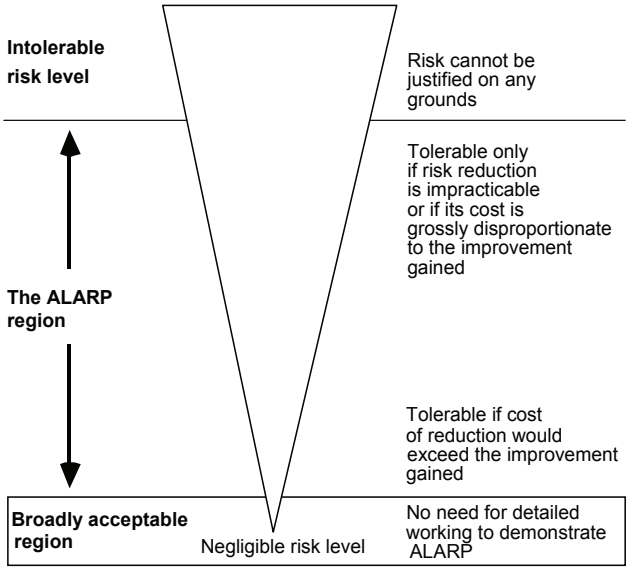

Another issue that can be raised in practice by the application of the ALARP principle is that one may have to be able confidently to evaluate these potential marginal variations before the detailed design and the implementation of the modifications are actually completed, or even started. Nevertheless, this does not rule out the application of the concept where risk reduction can only be judged qualitatively, rather than quantitatively. For example, the simple addition of a further safety feature, which costs relatively little, may be obviously worthwhile – qualitative judgements of this nature can often be readily made. Thus the application of ALARP should not be seen as restricted only to those circumstances which are amenable to quantitative reliability analysis.

2.2.4 Numerical targets and probabilistic safety analysis

A significant aspect of the nuclear safety approach is the role of numerical, probabilistic, targets. Numerical targets are explained in the SAPs[1] as:

> *'The individual risk of death levels in R2P21 cover risks to workers and to members of the public from activities on the site, which are:*
>
> *Boundary between the 'tolerable' and 'unacceptable' regions for risk entailing fatality:*
>
> *Worker: 1 in 1000 pa*
>
> *Member of the public: 1 in 10,000 pa*
>
> *Boundary between the 'broadly acceptable' and 'tolerable' regions for risk entailing fatality:*
>
> *Worker: 1 in 1 000 000 pa*
>
> *Member of the public: 1 in 1,000,000 pa'.*

The Safety of Operational Computer Systems (SOCS) report[22] continues with:

> *'Safety systems will feature amongst the risk reducing provisions comprised in this demonstration, which will thus include qualitative substantiations of compliance with appropriate safety engineering standards supplemented (where practicable) by probabilistic analyses of their reliabilities. Other techniques which may be used for structuring the safety case include fault and event tree analysis, failure mode and effec ts analysis (FMEA) and hazard and operability studies (HAZOPS).'*

Probabilistic safety analysis (PSA) is now an accepted aspect of the demonstration of safety. The PSA is based on the overall plant design and operation and covers all initiating events that are in the design bases. It is performed using best-estimate methods and data to demonstrate the acceptability of the plant risk. The PSA provides a comprehensive logical analysis of the plant and the roles played by the design for safety, the engineered safety systems and the operating procedures. The PSA also demonstrates that a balanced design* has been achieved. This means that no particular class of accident of the plant makes a disproportionate contribution to the overall risk, eg of the order of one tenth or greater. The PSA provides information on the reliability, maintenance and testing requirements for the safety and safety-related systems. The techniques used for safety analysis are various, with fault trees, event trees and FMEA being the dominant methods, and HAZOPs being used in fuel reprocessing applications.

* Balance in design refers to the concept that one feature should not be used disproportionately to another or that the integrity requirements on a feature are commensurate with its actual integrity

The PSA provides a check on the level of safety achieved in terms of plant risk and provides a means of claiming that the risk is lower than the intolerable region established by the NII's SAPs.[1] The PSA can then be used to demonstrate that risks are ALARP by investigating the effect on plant risk of modifying the plant safety provisions. It follows that levels of unreliability will have been ascribed to the systems that determine the overall plant risk, and the achievement of these levels will need to be demonstrated in a robust manner in the safety case. However, there are some aspects of safety that the principles recognise as not readily amenable to simple quantification of failure. The role of human factors (at an individual, group and organisational level) in achieving safety and initiating accidents is hard to quantify meaningfully, especially where knowledge-based activities are concerned. Similarly, the contribution of good management practices is hard to assess, although research and some progress has been made in this area. Other areas identified that are difficult to quantify are common mode failures and other types of failure due to design faults or specification omissions. The latter are particularly important from the point of view of this report, since they include failures due to software faults.

For many, the potential for societal consequences of a severe nuclear accident suggests that the requirements for nuclear reactor protection systems are very onerous. However, from the point of view of risk-based criteria, the use of diversity, defence in depth and the infrequent demands on the systems mean that in the UK, protection systems have had requirements for a probability of failure on demand of 10-3 to 10-4. The protection systems for New Build may have more stringent targets but these have yet to be reconciled with the claim limits suggested by the IAEA. It is difficult to compare these figures directly with devices that have to continuously operate but it is generally accepted that these computer-based probabilistic requirements are modest when compared with, for example, avionics requirements. Nuclear protection systems have about one demand per year so a required probability of 10-4 failures per demand can be translated to about 10-8 failures per hour.

It is instructive to compare the nuclear requirements with those for a health application. One starting point might be to consider that an infusion pump manufacturer would not want a population of devices (say 100k operating at any one time), over a lifetime of the product (five years?) to have a critical safety incident. If we put the risk of this critical incident at 10-2 we get a required probability of failure of about 10-2 in 4 x 109 hrs or a rate of ~2 10-12 failures per hour. This a very small number, about three to four orders of magnitude lower than the requirements on a reactor protection system.

However, if we did a calculation of the comparative risks of a patient undergoing treatment and required that the risks from the equipment he/she is attached to should be very much less than any other risks, the requirements might not be quite so onerous. If we take the 'worker' figure of ~10-7 per hour and say that there may be 10-100 sources of harm, each of which should have a small (0.01) impact, we get a figure of 10-10-10-11 failures per hour. This is still an ambitiously small figure. If, on the other hand, the patient is at a very high risk from other sources, then the figure for tolerable risk would be much higher and the equipment figure scaled proportionally.

2.2.5 Objectives of the safety case

We have already stated the purposes of the safety case: to provide written substantiation of the plant safety, incorporating justification that standards are complied with and that risks are reduced to ALARP.

Given the magnitude and complexity of the legislative and technical requirements that have to be met, safety cases have to be structured in a logical manner and be demonstrably complete. The safety case has to support an argument that the requirements placed upon it are met. As such, the safety case contains *claims* about the properties of the system and, following a systematic approach, these claims are substantiated by evidence. Safety cases can be seen to support the following:[3]

– **Reasoning and argumentation**. A safety case can be seen as an overarching argumentation framework that allows us to reason as formally as necessary about all the claims being made. Here, there are two very different viewpoints: the first viewpoint is that the argument part of a safety case is mainly made by using prose which explains the connections between claims and evidence, and the second viewpoint is that the argument part of a safety case is made by

using a formal notation or structure which sets out the connections precisely. There are some hybrid approaches where the case can be seen to integrate and communicate a selection of formal analyses and evidence. For example, the formal reasoning about component timing could be recorded in a separate analysis.

– **Negotiation, communication, trust**. The safety case represents a boundary object between the different stakeholders who have to agree (or not) the claims being made about the system. To this end it has to be detailed and rigorous enough to effectively communicate the case and allow challenges and the subsequent deepening of the case.

The NII carries out periodic inspections at all nuclear sites to ensure that the licence conditions and regulatory requirements are all met. The safety case is central to these as it describes how the licensee meets these requirements and as such may facilitate inspections. An inspection would typically consider areas such as the following:[38]

– emergency preparedness
– incidents on the site
– staff training, qualifications and experience
– management of operations, including control and supervision
– examination, maintenance, inspection and testing
– organisational changes
– plant construction
– industrial safety.

To serve these purposes, the operator and the regulator interact regularly, ensuring that the safety case is current and acceptable. Both stakeholders have responsibilities throughout the plant's lifecycle which are stated by the Health and Safety Executive – these are discussed in the following section.

2.2.6 Safety case lifecycle and responsibilities

The following section describes these responsibilities and has been extracted from the TAG that details the 'Guidance on the Purpose, Scope and Content of Nuclear Safety Cases'.[6]

2.2.6.1 Production

The responsibilities for production, revision, review and document control should be clearly defined as part of licence compliance arrangements and be discharged by suitably qualified and experienced people. Where the licensee itself does not produce all of the safety case and uses contractors for this purpose, at all times the licensee must possess (in-house) the technical capability to understand its safety case and act as an 'intelligent customer'.

The responsibility for producing and maintaining a safety case may change as the plant moves through its lifecycle. For example, in the design stage, the safety case may be developed and owned by a design team who eventually hand over responsibility and ownership to the 'operator'. The safety management system should explain how relevant information is transferred (eg there should be a system in place to take forward design assumptions into operations) and demonstrate that there are mechanisms in place to ensure that the safety case is fully adopted and implemented. This also applies to a plant at the end of its operational life, when responsibility and ownership for the safety case may pass from the operator to a decommissioning group.

2.2.6.2 Peer review and independent assessment

As part of the production process, a safety case should undergo appropriate verification controls and a formal approval process to check, among other things, that: the safety case is complete; key safety assumptions are valid and have been subject to a sensitivity check; appropriate robust methods and data have been used; calculations are correct; and the plant and operational details documented are consistent with the actual plant and its operations. In addition, and where necessary, there should be independent safety assessment by suitably qualified and experienced assessors, who are independent of the authors and verifiers and those directly responsible for the plant's operations. Following independent assessment, safety cases should be considered by the licensee's Nuclear Safety Committee (NSC). When the licensee's arrangements contain a safety classification system based on safety significance, these arrangements may restrict the independent safety assessment and reference to the Nuclear Safety Committee to the more significant safety cases.

2.2.6.3 Maintenance

It is important that the safety case is kept up to date. Significant changes may occur during operations such as modification, incidents, revised life expectancy, etc. Such changes should be recorded and taken forward as necessary in an updated safety case, which accurately and readily reflects the current situation.

Where referenced data and information underpin analyses and assumptions, arrangements are needed to ensure that when relevant new information comes to light a review is conducted to check whether and how safety cases they support are affected. Documentation which no longer forms part of a current safety case, or which has been superseded, should be identified and archived.

2.2.6.4 Periodic review

The NIA 1965[4] requires that 'the licensee shall make and implement adequate arrangements for the periodic and systematic review and reassessment of safety cases'. The purpose of this licence condition is to ensure that throughout its life, each plant remains adequately safe and that its safety case is kept up to date.

Two types of reviews are required: interim reviews, and major safety reviews. The latter are commonly referred to as periodic safety reviews (PSRs).

Interim reviews are carried out to provide regular confirmation that the safety case remains valid and that the safety of mid-term future operations will continue to be demonstrated by the case. They should cater for components whose behaviour or nature may change significantly and, if necessary bring forward the date of the next PSR. Such reviews would normally be expected every one to three years (eg at the time of periodic outage for reactors). The licensee's arrangements should also initiate reviews if new information indicates any significant change in safety case assumptions.

The purpose of a PSR is to determine, by means of a comprehensive assessment, whether the plants, processes, management, operations and facilities covered by a safety case remain as safe as reasonably practicable when judged against modern standards. It should also determine that ageing and other time-related phenomena will not compromise safety, particularly before the next PSR. The maximum period between PSRs is normally ten years.

2.3 Discussion

This section briefly presented the regulatory framework within which nuclear safety is assured in the UK. A licence to operate a nuclear installation is required by the Health and Safety at Work etc Act 1974 and the Nuclear Installations Act 1965. The legislation places the HSE responsible for granting that licence. The Nuclear Installations Inspectorate is responsible for reviewing and approving that licence. This is done on the basis of the safety case; the safety case has to demonstrate that risks have been reduced to ALARP and that all applicable legislation and regulations have been met.

These requirements are explained in more detail in the SAPs and the TAGs, developed and maintained by the HSE and NII. These publications state the principles for assessing the safety case, for which the operator is entirely responsible. International standards address in more detail technical aspects regarding the engineering of systems in nuclear installations.

The safety case is a living entity that is maintained by the operator, with support from third parties, such as suppliers and consultants, and is periodically reviewed by the NII. The operator and the regulator interface to facilitate the licensing process. Although the operator is responsible for the safety case, the regulator may give advice to the operator as to how the requirements can be met.

The UK regulatory framework around nuclear safety has been shaped over the past 60 years or so. There have been significant events in this recent history that are important in understanding the significance of certain key aspects of the regulation. The UK and international nuclear domain have learnt grave lessons from accidents that have occurred in the recent past. These, along with other economical, technological and societal conditions, have had a great impact on safety regulation and have made the safety case an important part of nuclear safety assurance.

3 Development and drivers

This section presents an overview of the development of the nuclear sector from 1950 onwards and discusses how various events, including accidents, competition with other energy sources such as gas, and the privatisation of the UK energy sector, have shaped safety regulation in the UK and elsewhere.

3.1 Chronology

The beginning – 1950

The UK's civil nuclear power programme grew out of the post-war military imperative of producing plutonium for nuclear weapons. The first two reactors were built in 1950, known as the Windscale piles at Sellafield, Cumbria. These only produced plutonium; however, a new design, where heat from the reactors was used to generate steam, which was in turn utilised to generate electricity, was developed. This new design was implemented with Calder Hall, the UK's first commercial nuclear power reactor, which began operating in 1956[9].

The first authority to regulate atomic energy and radioactive substances, the UK Atomic Energy Authority (UKAEA), was established in 1954.

A programme of nuclear power – 1955

The government of the time recognised the importance of providing the nation with alternative energy sources due to the foreseen shortage of fossil fuel. In addition, the Calder Hall plant – which had been put in operation only a mere 14 years since the initiation of the first, experimental reactor in history – was a great success and was seen as a driver to expand the UK nuclear programme. In 1955, the government published a white paper titled 'A programme of Nuclear Power', which outlined a ten-year plan for the construction of 12 nuclear power stations. The paper showed confidence that there was little possibility of environmental contamination and opted for a further large-scale expansion of nuclear power after 1965.

The Windscale fire – 1957

However, in 1957, an accident occurred. On 8 October 1957, the graphite core of a nuclear reactor at Windscale, Cumberland (now Sellafield, Cumbria) caught fire, releasing substantial amounts of radioactive contamination into the surrounding area. The event, known as the Windscale fire, is the worst nuclear accident in Great Britain. It is noteworthy that the fire was ongoing for about 48 hours before it was eventually detected. The responders were uncertain as to how to control it and made several failed attempts to put it out. Eventually, it was decided to pour water directly on the fire, which involved a high risk of explosion (molten metal oxidises in contact with water, stripping oxygen from the water molecules and leaving free hydrogen, which could mix with incoming air and explode). This attempt was successful although eventually the impact on human life and the environment was dramatic. It has been estimated that the incident caused 240 cancer cases.[10]

The aftermath of the Windscale fire: the Nuclear Installations Act 1959 and the introduction of the Nuclear Installations Inspectorate

Various committees were appointed to carry out public inquiries into the causes of the fire. The inquiries looked into, not only the events of the accident timeframe, but also organisational and policy issues that allow for the failure to occur. These inquiries (and in particular the Fleck inquiry) had a significant impact on the structure and function of the UKAEA. One of the results of the inquiry was to put in place, within the UKAEA, an Authority Health and Safety Branch (AHSB), to separate safety supervision from plant operation.[11]

Another result of the Windscale fire inquiries was the establishment of the Nuclear Installations Act of July 1959. This required that the civil nuclear power stations which were under construction and those planned for the future, which were to be operated by the Central Electricity Generating Board (CEGB), would be licensed by the newly formed Nuclear Installations Inspectorate (NII).

It can be seen that the failures which led to the Windscale fire and its poor handling underlined the need to have in place a powerful regulator whose only concern was the safety of the nuclear power programme. However, the fire did not have an impact on the UK nuclear power station construction programme, which proceeded without a pause.

Expansion of the UK's nuclear programme – 1957-1971

In a period of 14 years (1957-1971), eight new, more advanced and larger power plants were built and commissioned, as planned in the 'A Programme of Nuclear Power' white paper of 1955. In 1964, a second white paper was published by the government which outlined a second nuclear power programme with four new stations for England and Wales to be in operation between 1970 and 1975.

In 1965, the Nuclear Installations Act that was initially released in 1959 was revised. The Nuclear Installations Act 1965 is still the main nuclear legislative publication.[4] (See Section 2.1.1.2 for more).

A third UK nuclear programme and developments elsewhere – 1970-1979

In these years, several other countries were also developing and expanding programmes of nuclear power. In 1973, the French government stated their intention to build 'not more than six or seven stations per year'. In the same year in the USA, the oil crisis gave impetus to the construction of several nuclear plants. Also in 1973, the UK's Central Electricity Generating Board (CEGB) stated its intention to develop 18 large nuclear power stations. The rapid expansion of nuclear programmes was attributed to a combination of factors: the massive increase of electricity demand between 1955 and 1975, concerns about shortage of fuels, and significant technological advancements in nuclear generation.

Three Mile Island accident – 1979

On 28 March, 1979, there was a partial core meltdown in Unit 2 of the Three Mile Island Nuclear Generating Station in Dauphin County, Pennsylvania, USA. The accident was caused by some mechanical failures; in addition, there were several human factors issues that significantly contributed to the incident and its magnitude. The user interface of the control system was poorly designed, confusing the operators, not allowing them to diagnose the problem for several hours.[12]

Although the radiological impact was negligible (the USA government inquiry concluded that no cancers were a result of the accident and there was very small environmental damage, although this has been debated by researchers), there were several psychological, financial and administrative consequences. It is notable that since the Three Mile Island accident, the US Nuclear Regulatory Commission has not reviewed an application to build a new nuclear power plant.

Another significant impact of the Three Mile Island disaster was the changes it brought about in safety engineering. In particular, since then, there has been increasing attention both in academic research and industrial practice to human factors engineering, emergency response planning, reactor operator training, radiation protection, and many other areas of nuclear power plant operations.

The anti-nuclear movement and the Sizewell Inquiry – 1980-1986

During the years shortly after the Three Mile Island disaster, there were several large anti-nuclear demonstrations. Concerns about the environment and public safety led to protests from 1979-1983, most notably in the USA and Germany, involving 100,000–600,000 people.

At the time in the UK, two advanced gas-cooled reactor (AGR) stations were being built (Heysham II and Torness), each having two reactors. There were significant challenges in the construction of the AGRs, which led to them falling behind schedule by several years. The problems had to do with the aerodynamic design of the fuel and the fact that most of the assembly and high technology construction had to be done on site. Taking the high costs and technical problems in mind while looking at the future, in 1980, the CEGB decided to use an American design – the Pressurized Water Reactor (PWR) – which would solve several of the technical and financial challenges of AGRs.

However, the PWR was, for the UK, a new type of civil reactor and in 1983, it was decided that a public inquiry would be held to review the PWR's acceptability. The public inquiry into Sizewell B's PWR was the longest-running inquiry in Britain's history, commencing in June 1983, and closing in March 1985. The review was based on the pre-construction safety case, which was submitted to the NII in August 1981. Sir Frank Layfield, who chaired the inquiry committee, reported in early 1987 that the safety case was satisfactory and the NII issued a licence to proceed with construction in August 1987. Despite the initial controversy that surrounded the original PWR decision, the design is now considered to be the safest design and is the most widespread around the world.[14]

The Chernobyl accident – 1986

The Chernobyl disaster is the worst civil nuclear power plant accident in history. According to an official report compiled on behalf of the Chernobyl forum,[15] the reactor vessel rupture and the series of explosions that followed in 26 April 1986 resulted in the deaths, within a few days or weeks. of 30 power plant employees and firemen (including 28 deaths that were due to radiation exposure), and it also brought about the evacuation of about 116,000 people from areas surrounding the reactor during 1986. According to the report, the number of cancer-related casualties could reach 4,000. The radioactive fall-out into the atmosphere drifted over large parts of the western Soviet Union and Eastern, Western and Northern Europe, with significant, long-term environmental impacts.

The Chernobyl accident raised concerns about nuclear safety internationally and resulted in a significant slowdown in the expansion of nuclear programmes around the world. Shortly after the accident, the World Association of Nuclear Operators (WANO) was established. WANO is an international non-profit organisation, based in the UK, with offices in Russia, the USA, Japan and France, with approximately 130 members from 30 countries. The objective of WANO is to help its members (nuclear operators) achieve high levels of safety and reliability. They offer peer reviews and technical support and have established a forum for operators to communicate lessons learned. In addition, the International Atomic Energy Agency (IAEA) carried out safety review projects for each particular type of Soviet reactor.[16]

Furthermore, since the Chernobyl disaster, much discussion has focused on the notion of safety culture, as defined by the IAEA: 'the personal dedication and accountability of all individuals engaged in any activity which has a bearing on the safety of nuclear power plants'. It has become widely understood that nuclear safety is not only a matter of engineering excellence, but a day-to-day concern of everyone involved.

Privatisation of the UK energy market – 1990-1995

The Thatcher administration saw, in privatisation of the energy sector, the potential for increasing efficiency, delivering lower prices for the consumers and offering more diversity. Firstly, the energy market was broken up in 1990 across the UK. The CEGB, the Scottish and Welsh systems were all broken up in multiple entities. Although most parts of the electricity generation and distribution systems were privatised, the nuclear elements were not; this was because of the completion costs of Sizewell B, which was still being built, and poor performance history.

In addition, gas was introduced as a source of energy in the early 1990s. By the mid-1990s the UK's North Sea gas system was linked up to a European gas network. Despite the initial benefits of gas energy, it was seen that the UK reserves were limited, thus having gas prices dominated by the international market, especially as dependency on gas increased.

The UK nuclear fleet, Sizewell B and British Energy – 1990-1996

In the meantime, the cost performance of the UK's AGR fleet was improving under Nuclear Electric plc and Scottish nuclear. It is also stated that there were significant improvements in safety performance. In addition, Sizewell B was finally completed in 1996. These developments led to the establishment of British Energy (BE), a new privatised nuclear entity.

Turbulence in the energy market – 1996-2008

As electricity prices rose at the end of the 1990s, gas remained cheap. By 1999, gas capacity had grown to be about 40% greater than nuclear.[37] As a result, and in combination with other organisational issues, British Energy suffered significant losses, from the beginning of its privatisation, eventually reaching financial restructuring. For the following years, BE underwent several structural changes, being unable to maintain a stable management structure (in 2004 it was restructured with UK government investment of more than £3 billion, although this has since been paid back in full).

However, in the years 2002-2006, crude oil prices saw a dramatic rise (from $22.8 to $65). These massive rises in oil prices had a cascade effect on European wholesale gas prices and in contract prices for UK wholesale gas, which also rose significantly, feeding through into the UK electricity prices. By 2006, British Energy was reformed, with greater managerial stability; in combination with the price rises in oil and gas, BE's fortunes improved considerably. In addition,

Government energy reviews over these years concluded that there should be an increase in the proportion of renewable energy in the electricity sector. As a result of these events, BE became more viable commercially.

The future: new build – 2008-2025

In September 2008, the acquisition of British Energy by Électricité de France (EDF) was agreed.[20] The acquisition took place in January 2009, when British Energy was bought for approximately £12 billion. Originally in 2008, EDF had proposed that it would build four European pressurised reactors (EPRs), each costing up to £2 billion, at two of British Energy's eight nuclear sites: Hinkley Point in Somerset and Sizewell in Suffolk.

However, in 2009, a wider expansion of nuclear power was signalled by the UK government as it named 10 sites where new power stations could be built. The first is set to be operational by 2018 and, by 2025, nuclear electricity generation could amount to around 40% of new energy provision.

Since the decision to build these new plans, the NII has revised the safety assessment principles (SAPs) and technical assessment guides (TAGs). This is because the SAPs and TAGs had been developed primarly considering the safety of existing plants, but not new ones (the last plant to be commissioned was Sizewell B in 1995).

Currently, the HSE is assessing two new power station designs:

- the UK-EPR designed by Areva and EDF
- the AP1000 designed by Westinghouse.

More background on the assessment process and the power station designs can be found in the HSE's report, *New Reactors*.[21] The assessment process that HSE is taking has the following characteristics.

- The regulators get involved with designers from the earliest possible stage, where they have most influence.
- Assessments are following a step-wise process, whereby they are getting increasingly detailed. This allows the regulators to identify issues early in the process and reduce the financial and regulatory risks for potential operators.
- Design issues are separated from specific site-related issues, improving the overall efficiency of the regulatory process.
- It is open and transparent. Detailed design information is available on the Internet, and the public is welcome to comment on it. The regulators also provide regular feedback on the progress of the assessments.

Apart from the assessment of the designs, the licensing requirements will still be placed on the approved designs. This means that the approval of the designs is only the first step; detailed safety cases will have to be produced and will be reviewed by the regulator.

3.2 Goal-based safety cases

The nuclear industry was part of a widespread adoption of such regimes and there are more general drivers to the approach than just the nuclear-specific factors mentioned above. The Robens Report[31] and the Cullen Inquiry[30] into the Piper Alpha disaster were major drivers behind the UK regulatory agencies introducing goal-based regulations. The reports noted several shortcomings with prescriptive safety regulations – that is, regulations that provide a strict definition of how to achieve the desired outcome.

Firstly, with prescriptive regulations, the service provider is only required to carry out the mandated actions to discharge its legal responsibilities. If these actions then prove to be insufficient to prevent a subsequent accident, it is the regulations and those that set them that are seen to be deficient. Thus safety can be viewed as the responsibility of the regulator and not the service provider whose responsibility, in law, it actually is.

Secondly, prescriptive regulations tend to be a distillation of past experience and, as such, can prove to be inappropriate or, at worst, create unnecessary dangers in industries that are technically innovative.

Thirdly, prescriptive regulations encode the best engineering practice at the time that they were written and rapidly become deficient where best practice is changing, eg with evolving technologies. In fact, it is quite probable that prescriptive regulations eventually prevent the service provider from adopting current best practice.

Another driver for adopting goal-based regulation, from a legal viewpoint, is that overly-restrictive regulation may be viewed as a barrier to open markets. Various international agreements, EC directives and regulations are intended to promote open markets and equivalent safety across nations. While it is necessary to prescribe interoperability requirements and minimum levels of safety, prescription in other areas would defeat the aim of facilitating open markets and competition.

Finally, from a commercial viewpoint, prescriptive regulations could affect the cost and technical quality of available solutions provided by commercial suppliers. So there can be clear benefits in adopting a goal-based approach, as it gives greater freedom in developing technical solutions and accommodating different standards.

A system safety case is now a requirement in many safety standards and regulations. Explicit safety cases are required for military systems, the offshore oil industry, rail transport and the nuclear industry. For example, in the UK, a nuclear safety case must demonstrate, by one or other means, the achievement of ALARP. In the Health and Safety Commission's submission to the government's 'Nuclear Review',† a safety case is defined as 'a suite of documents providing a written demonstration that risks have been reduced to ALARP. It is intended to be a living dossier which underpins every safety-related decision made by the licensee.'

The system safety case, of course, varies from sector to sector. The core of a nuclear system safety case is (i) a deterministic analysis of the hazards and faults which could arise and cause injury, disability or loss of life from the plant either on or off the site, and (ii) a demonstration of the sufficiency and adequacy of the provisions (engineering and procedural) for ensuring that the combined frequencies of such events will be acceptably low. Safety systems will feature among the risk-reducing provisions comprised in this demonstration, which will thus include qualitative substantiations of compliance with appropriate safety engineering standards supplemented (where practicable) by probabilistic analyses of their reliabilities. Other techniques which may be used for structuring the safety case include fault and event tree analysis (FTA and ETA), failure mode and effects analysis (FMEA) and hazard and operability studies (HAZOPS).

The safety case, particularly for computer-based systems, traditionally contains diverse arguments that support its claims. These arguments are sometimes called the 'legs' of the safety case and are based on different evidence. Just as there is defence in depth in employing diversity at system architecture level, so we see an analogous approach within the safety case itself. Another important feature of the safety case process is independent assessment. The objective of independent assessment is to ensure that more than one person or team sees the evidence so as to overcome possible conflicts of interest and blinkered views that may arise from a single assessment. The existence of an independent assessor can also motivate the assessed organisation. The relationship between independent assessment and 'legs' can, however, be complex.

Safety cases are important not only to minimise safety risks but also to reduce commercial and project risks. In industries such as the nuclear industry, the need to demonstrate safety to a regulator can be a major commercial risk.

To sum up, the motivation for a safety case is to:

- provide an assurance viewpoint that demonstrates that safety properties are satisfied and risks have been satisfactorily mitigated
- provide a mechanism for efficient review and the involvement of all stakeholders
- provide a focus and rationale for safety activities
- demonstrate discharge duty to public and shareholders
- allow interworking between different standards and support innovation.

Therefore, in a safety case, the emphasis should be on the behaviour of product, not just the process used to develop it. A useful slogan is 'What has been achieved, not how hard you have tried'.

† The review of the future of nuclear power in the UK's electricity supply industry.

3.3 Discussion

This section presented the chronology of events that shaped the UK nuclear industry and its safety regulation.

In brief, in addition to the cross-sector drivers for goal-based approaches, we see the following drivers of nuclear safety and the safety case.

3.3.1 Major disasters

The three major accidents discussed here, apart from the significant human, environmental, social and economic consequences, also each eventually resulted in a series of legislative, policy and engineering changes following the subsequent investigations. The Nuclear Installations Act 1959 (and then the revised release of 1965) – a direct result of the Windscale fire – set the requirement for licensing of the nuclear site and the establishment of a nuclear regulator (the NII). Since the Windscale fire, there has not been another nuclear accident in the UK.

However, the Three Mile Island accident, and more dramatically, the Chernobyl disaster, had a massive impact on the nuclear sector internationally; they resulted in a slowdown of nuclear programmes across the world, while the findings of the subsequent investigations into the causation of these accidents helped increase the focus on safety and reliability of nuclear power plants, systematically considering aspects such as human factors, operator training and protection systems, and establishing a safety culture.

3.3.2 Market turbulence

Privatisation of the energy sector, the introduction of gas and raises in oil prices had a significant impact on the nuclear sector over the past 20 years. Although the nuclear sector went through significant challenges during privatisation of the energy market and competition with gas in the 1990s, British Energy eventually reached a successful level before being acquired by EDF. This led to the decision to proceed with the development of 10 new power plants in the UK.

3.3.2 Technological innovation and the importance of the safety case

In the UK, the Sizewell B inquiry identified the importance of the nuclear safety case; the pre-construction safety case for the PWR design underwent a two-year long review by the NII, with the final conclusion being favourable, and despite the significant amount of controversy, proving to be the most 'inherently' safe design to date. Sizewell B was the last nuclear power plant to be commissioned, in 1995.

Similarly, the decision to proceed with new build has had an important impact on the SAPs, while the regulator has established a careful review process to approve the designs. Once the designs are approved, the licensing process will commence and again will be based on the safety cases that will have to be developed by the operator.

4 Approaches to safety cases and content

In this section, we focus on safety case practice; we look at approaches to the development of safety cases and take a closer look at the contents of a safety case based on the HSE SAPs.

4.1 Approaches to safety cases

There are different strategies that can be deployed in developing the safety justification. The three main approaches can be characterised as a 'triangle' of the use of accepted standards and guidelines, justification via a set of claims/goals about the system's safety behaviour, and an investigation of known potential vulnerabilities of the system.

This is illustrated in Figure D3 below.

Figure D3: The 'triangle' of safety justification

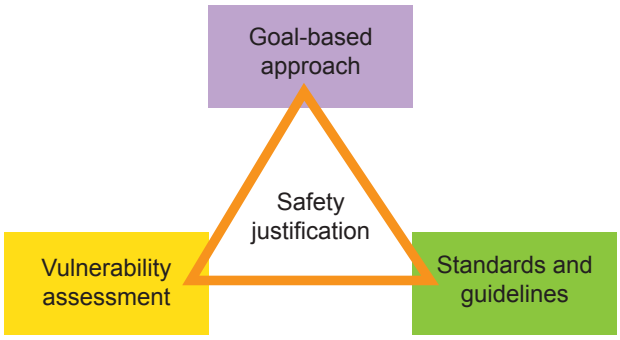

4.1.1 Standards compliance

The first approach is based on demonstrating compliance to a known safety standard. This is a common strategy.

Demonstration of compliance is not always a straightforward task. In the nuclear sector for instance, for a particular type of device – smart instruments – an approach has been developed to assess their development process against IEC 61508.[33]

4.1.2 Goal-based approach

The second approach is goal-based (as discussed in Section 3.2) – where specific safety goals for the systems are supported by arguments and evidence at progressively more detailed levels.

A goal-based safety case typically takes the form of claims-arguments-evidence. A claim is made about a property of the system (eg the system is acceptably safe) and the evidence is provided to support this claim. The argument has the purpose of explaining the approach undertaken to support the claim (eg 'argue by considering safety of sub-systems').

The structured safety case needs to be challenged and assessed for it to be considered to be fit for purpose. In some areas, such as defence and nuclear, there is a well-defined process for such independent assessment.

The basic measure of efficacy of an argument in this work is the *confidence* that the argument engenders in a dependability claim. Informally here, a safety case is taken to be some *reasoning*, based upon *assumptions* and *evidence*, allowing certain *confidence* to be placed in a dependability *claim*. For a given claim (eg *pfd* is smaller than 10-3), the confidence – and its complement, *doubt* – will depend upon confidence/doubt in the truth of assumptions, in correctness of reasoning, and in 'strength' of evidence.

4.1.3 Vulnerability-based approach

The final approach is a vulnerability-based argument, where it is demonstrated that potential vulnerabilities within a system do not constitute a problem. This is essentially a 'bottom-up' approach as opposed to the 'top-down' approach used in goal-based methods.

These approaches are not mutually exclusive, and a combination can be used to support a safety justification, especially where the system consists of both off-the-shelf components and application-specific elements.

4.2 Contents of a safety case (SAPs)

The SAPs[1] detail the information that the safety case should contain in order to demonstrate safety. This information should be easily accessible and understandable.

According to the SAPs, '…answers to the following questions should be obtained:

- What is the safety case for (a new site/facility, plant extension, modification)?
- What does the site/plant, etc. look like (site layout, design, key features)?
- What must be right and why (e.g. structural integrity, performance)?
- How is this achieved (e.g. codes, standards, specifications)?
- What can go wrong (faults, hazards)?
- What prevents/mitigates against it going wrong (e.g. protection systems, redundancy, diversity, procedures)?
- What if it still goes wrong (risk/consequences, emergency arrangements)?
- What could be done to make it safer ('Optioneering' and ALARP considerations)?
- What must be done to implement the safety case (e.g. operating procedures, limits and conditions, maintenance)?
- How long will the safety case be valid (e.g. full life time or shorter due to life limiting features)?
- What happens at the end-of-life (decommissioning principles / strategy)?'

The SAPs continue to state that:

> *'A safety case may comprise a hierarchy of documents. The top tier will contain the core of the safety arguments and increasingly detailed technical documents and supporting analysis will be presented in lower tiers. At the lowest level there are likely to be the engineering calculations, experimental results and data on reliability and relevant operational experience, etc. At the other end of the scale, for example long-term storage of low level radioactive material, a relatively simple safety case is normally appropriate*

that sets out requirements for periodic inspection to ensure that containment remains sound and that storage conditions remain conducive to long-term stability.'

The SAPs also provide examples of the different types of documented information which underpin the safety arguments. These may include:

- identification of faults and hazards and compilation of a comprehensive fault schedule
- criteria for choosing the design basis faults and hazards (ie those faults and hazards for which design measures are explicitly claimed in the safety case).
- deterministic analysis of the design against these faults to show a robust tolerance of them
- determination of the safety functional requirements of the structures, systems and components important to safety
- determination of limits, conditions and associated trip and alarm settings and a comprehensive protection schedule
- task analysis of important operations
- substantiation that the plant will deliver the safety requirements
- probabilistic safety analysis
- identification of suitable emergency arrangements.

The precise structure and scope of the documentation will be a matter for the licensee to determine, taking into account the significance of the hazard and complexity of the safety case.

The tables below, taken from the TAGs,[6] identify core requirements and activities that will feed into a nuclear safety case.

Table D3: Sources of requirements

Sources of Safety requirements	Description
Safety policy	Principles and objectives
Safety criteria	Statutory limits
Safety standards	International, national and licensee-specific codes and standards
Research and development	To determine appropriate criteria, standards, etc.

Table D4: Typical approaches to safety demonstration

Approaches to safety demonstration (the safety case)	Description
Deterministic analysis	The identification of limits and conditions in the interest of safety
Engineering substantiation	To show safety function delivery with appropriate integrity
Probabilistic safety analysis	Including sensitivity analysis
ALARP arguments	Including options considered

Table D5: Maintenance of the safety case

Monitoring and maintenance of the safety case	Description
Safety case management process	Including modifications to plant and safety case
Periodic review	Including interim review
Operating data	Including incidents

Going in more depth, the SAPs[1] contain a number of engineering principles that are to be followed and demonstrated in the safety case. The principles in this section are presented in three main groups (key principles, general principles and engineering principles for specific areas).

A selection of engineering for specific areas is presented below. These areas are:

- design for reliability and reliability claims
- commissioning, maintenance, inspection and testing
- ageing and degradation
- external and internal hazards
- safety systems and control and instrumentation of safety-related systems
- human factors
- control of nuclear matter
- containment and ventilation
- reactor core.

As an example, Table D6 provides an overview of the expectations set out by the engineering principles that concern human factors.

Table D6: Human factors requirements

Requirement	Description
Integration with design, assessment and management	An approach should be taken to integrate human factors within the design, assessment and management of systems.
Allocation of safety actions	The allocation of actions between humans and technology should be determined and substantiated. Dependence on human action to maintain a safe state should be minimised.
Identification of actions impacting safety	The thorough identification of all human interactions, which, if not performed correctly, could affect the fulfilment of the safety functions.
Identification of administrative controls	The identification of controls used to remain within safe operation. These should be designed so that requirements for personnel action are clearly identified and unambiguous to those responsible for their implementation.
Task analysis	Analysis of the tasks important to safety to determine demands of personnel in terms of perception, decision making and action.
Workspaces	Workspaces in which operations and maintenance of the plant take place should be designed taking into account human perceptual and physical characteristics. The design of workspaces should also consider the impact of environmental factors.
User interfaces	User interfaces should be provided in support of effective monitoring and control of the plant during all plant states. The interfaces, which may comprise indications, controls, instrumentation and alarms, should be designed taking into account human psychological, cognitive and physical characteristics.
Personnel competence	There should be a systematic approach to the identification and delivery of personnel competence. Job design, training programme design and implementation and competence assessment should be parts of this systematic approach, which should be applied to all personnel involved in actions that could have an impact on safety – this also includes contractors.
Procedures	Documented procedures should be produced to support human performance during activities which may have an impact on safety.
Human reliability	Risk assessment should identify the possibilities for error during human action that could have an impact on safety. Human error could take place during operation or maintenance as well as post-fault responses and long-term recovery.

In the lifecycle of a plant, from conception through to decommissioning, there are various key stages which require special consideration. The safety case for each stage should demonstrate the safety of that stage before it commences and should be forward looking to subsequent stages.

4.3 Safety systems (C&I)

In the nuclear sector, a particular focus is given to control and instrumentation (C&I) systems. C&I systems support the monitoring and control of the reactors, as they measure and/or regulate process variables such as flow, temperature, level, or pressure and also perform functions of the highest safety significance, such as emergency reactor shutdown.

Apart from their safety significance, in this report, we have decided to focus on C&I systems for the following reasons, which we also see relevant to healthcare.

- C&I systems may be computer controlled, thus containing software. This poses significant challenges to the safety justification and eventually the safety case.
- C&I systems may comprise of commercial off-the-shelf (COTS) components, such as programmable alarms.

In this section we will therefore provide more detail on the issues regarding safety justification of C&I systems.

4.4 Safety cases for systems containing software

Below is an extract from the Safety of Operational Computer Systems (SOCS) report,[22] written for the HSE.

> *Although computer technology changes rapidly in developing increased performance and penetrating new applications, it is hardly a new technology anymore. It would seem a reasonable aim that the production of software safety cases should become much more part of the overall engineering approach and not something that gives rise to undue project or technical risk. However, this is not the case at the moment, despite the considerable experience in both the nuclear and non-nuclear sectors in developing what are in effect software safety cases. The problems of software engineering in general are discussed elsewhere in this report. In terms of the safety case, past problems have been due to:*
>
> – *lack of initial explicit definition of the technical approach to be taken in justifying safety cases*
>
> – *confusion over what is necessary to achieve dependable systems and the evidence to justify a system: for example, many standards contain an undifferentiated mixture of project development and assurance activities*
>
> – *lack of a documented, agreed, accessible engineering process for computer safety cases (rather than for their development)*
>
> – *difficulty in justifying 'reasonably practicable' and demonstrating ALARP.*

For these reasons, a different approach is taken for systems containing software. In particular, the HSE SAPs require that:

> *'Where the system reliability is significantly dependent upon the performance of computer software, the establishment of and compliance with appropriate standards and practices throughout the software development life-cycle should be made, commensurate with the level of reliability required, by a demonstration of "production excellence" and "confidence-building" measures.'*

The two requirements are described as follows.

- 'Production excellence': The development lifecycle meets the requirements placed by all the applicable standards.
- 'Independent confidence-building measures': A combination of goal-based assessments (eg demonstration of accuracy, reliability, etc.) including assessment of potential vulnerabilities in the implementation of the system/component.

This approach is widely known in the nuclear safety domain as the 'two-legged approach' and is discussed below in more detail. There are also claim limits that apply to the system (see Appendix D1 for examples of this concept).

4.4.1 Production excellence

According to the SAPs,[1] it must be demonstrated that there is:

- thorough application of technical design practice consistent with current accepted standards for the development of software for computer-based safety systems
- implementation of an adequate quality assurance programme and plan in accordance with appropriate quality assurance standards
- application of a comprehensive testing programme formulated to check every system function, including:
  - prior to installation on site, the verification of all phases of the system production process and the validation of the integrated system against its requirements specification by persons not involved in the specification and design activities
  - following installation on site, a demonstration that the safety system, in conjunction with the plant, performs to requirements, this demonstration being devised by persons other than the system specifiers, designers or manufacturers
  - a programme of dynamic testing, applied to the complete system that is capable of demonstrating that the system meets its reliability requirements.

This evidence-gathering activity would usually be supported by an audit.

If weaknesses are found in the production process, compensating measures must be applied which are targeted at the specific weaknesses. These may be designed to rectify or mitigate the deficiencies.

4.4.2 Independent confidence-building measures

Independent confidence-building measures (ICBMs) are performed by the licensee or an independent agent contracted by the licensee. They should provide a thorough and challenging assessment of fitness for purpose, but should be reasonably practicable.

This will involve:

- complete and preferably diverse checking of the software (after validation has been completed) by a team independent of the suppliers
- independent assessment of the full test programme, covering the full scope of the testing activities.

Techniques employed for the product analysis aspect of independent confidence building might include

- static analysis of software source code
- dynamic analysis of software with test coverage metrics
- a post-installation demonstration, performed independently of the suppliers
- that the system performs to requirements
- analysis of data from previous use of the system in the nuclear industry or elsewhere.

For pre-developed components, additional techniques that may be part of the confidence-building programme include:

- independent review of supplier procedures and standards
- review of certifications awarded to the component or supplier
- review of operating data from prior uses of the component, provided they are obtained from reputable sources and relevant to the version of the component under consideration
- review of operating data from prior uses of different versions of the component, as general support for manufacturer excellence
- review of the supplier's track record as a producer of high-quality components for safety applications.

4.5 Current developments

The approach taken in the UK depends on the history of the system, the nuclear organisation involved, the category and role of the devices. Much work in the UK is refurbishment and enhancements – eg adding protection channels, replacing obsolete sensors, refurbishing old control systems on AGRs, and replacing UPSs (uninterruptible power supply).

There have been some significant developments since the commissioning of Sizewell B in 1995. An IAEA paper[27] from British Energy outlines some of the experience and emerging trends which are worth discussing here. Taking a particular focus on C&I, research overseen by the Control and Instrumentation (C&I) Nuclear Industry Forum (CINIF) has been on refurbishment, particularly addressing issues of the use of smart sensors. It mentions research on:

- diversity
- formal/static analysis of legacy code
- relative contributions of static and dynamic analysis
- statistical software testing
- low SIL applications[28]: PLCs, PCs
- safety justification of smart instruments and goal-based assessment of COTS products
- collaboration towards European harmonisation.

4.6 Summary

This section briefly presented the approaches to safety case development and the contents of a nuclear safety case. Three approaches are primarily identified: demonstration of compliance with standards, goal-based, and vulnerability-based.

The contents of a nuclear safety case are discussed in the SAPs and TAGs. In this section, we have paid particular attention to control and instrumentation (C&I) systems, due to their safety significance and the fact that they may contain software and COTS components. For these, the nuclear sector takes a particular approach not seen in other industries involving the evaluation of production excellence and independent confidence-building measures.

5 Lessons and recommendations for healthcare

Safety cases have a long history in the nuclear industry. Safety is achieved by addressing the complete landscape from the cultural to the technical, for each phase of the lifecycle, including formalised approaches for learning from experience. The industry approach to safety and the regulatory environment has evolved in response to technology drivers, the energy market, nuclear accidents and the need for an effective dialogue with a range of stakeholders. Safety cases are not something that is just imposed on the industry by the regulator, but are a shared approach to how safety should be justified and be seen to be justified. The safety case regime is accepted in the nuclear industry; the remaining issues are around how to make effective and efficient cases.

In terms of approaches to assessing computer-based systems, there are a number of aspects of the nuclear approach which are relevant to the health sector.

- Safety cases have an important role in both communicating and reasoning about safety and can also be used for the regulator as a basis for carrying out physical inspections.
- Safety assessment principles which guide the regulator are important to provide technical depth to overall goals of regulation, but should not prescribe solutions.
- Independent assessment and challenge are important parts of gaining confidence in systems – the nuclear sector uses independent confidence-building measures (ICBMs).
- Safety needs to be informed by security; in the nuclear sector, the Office for Civil Nuclear Security (OCNS) is embedded within the Nuclear Directorate.
- The nuclear safety case is applied within a context of high-risk hazards, public attitude to risks, licensed operators, trained users, and an international market.

- The industry, in the past, has made relatively modest reliability claims about protection system reliability, with claims above 10-4 probability of failure on demand (PFD) being rare. Claim limits on reliability figures are an important concept. Our approximate calculations suggest that the requirements on medical systems might well be more stringent than those on reactor protection systems.
- Fail-safety, dynamic design, the single-failure criterion, defence in depth, and diversity are important design principles.
- Standards and guidelines are important parts of the excellence of production argument.
- The classification of systems into categories of different safety significance is an important part of the safety case and overall plant design.
- The assessment of computer-based systems is embedded in the overall system and plant engineering processes. Safety justification requires both deterministic arguments and probabilistic safety analysis (PSA).
- There is considerable long-term research on goal-based structures for safety cases and different approaches for justifying systems. The healthcare industry should note that the safety case approach is more than just the structured cases presented using claims-argument-evidence (CAE) notation or goal structuring notation (GSN).
- A variety of different static analysis techniques and statistical testing is used in the assessment of computer-based systems.
- The industry is conservative and the rate of change is slow. The technology adoption trajectory involves extensive research, case studies, field trials, initial deployment, evaluation, adaptation and improvement.

In terms of challenges for healthcare, it might be useful to structure the issues around the following ideas.

- *Culture and technology maturity*. The safety culture prerequisites required for a safety case approach and how they can be identified and assessed. How the extent of changes needed can be described and a healthcare-specific approach developed.
- *Communication*. The challenge of developing an approach to justifying devices: separating issues of risk communication from problems of evaluation.
- *Evaluation and evidence generation*. The need to improve the evidence base and technologies used to justify a device. Cases might initially lead to clearer assessment of product weaknesses and strengths and highlight areas of debate.
- *Challenge and confidence building*. How cases are progressively challenged and developed into something that all stakeholders are confident in. There are issues of rigour and stopping rules and what is appropriate in different healthcare applications.

Finally, in relation to these we could consider how the CAE approach may be applied in healthcare.

6 Conclusions

This report explains the UK nuclear safety regulatory approach and the role of the safety case in the licensing process. The safety principles of the UK regulatory assessment are now well established and clearly articulated in the SAPs and will be applied to new build as well as the refurbishing of existing stations. The onus on the licensee to make a case and the overall goal-based approach to regulation in the UK provides for flexibility in the licensing. This flexibility also introduces some uncertainties and project risks.

There is currently much activity in the development of approaches within the overall frameworks provided by the SAPs and international standards. The use of ALARP, claim limits and diversity, excellence of production and independent confidence-building and their assessments are key parts of the UK approach.

7 Glossary

Term / Abbreviation	Explanation
ACSNI	Advisory Committee on the Safety of Nuclear Installations
BE	British Energy
CAE	claims-arguments-evidence
CEGB	Central Electricity Generating Board
EA	Environment Agency
GSN	goal structuring notation
HSE	Health and Safety Executive
HSW	Health and Safety at Work Act, 1974
IAEA	International Atomic Energy Agency
IRRs	Ionising Radiations Regulations 1999
MoD	Ministry of Defence
ND	Nuclear Directorate of HSE
NIA	Nuclear Installations Act 1965
NII	Nuclear Installations Inspectorate
NSC	Nuclear Safety Committee
OCNS	Office for Civil Nuclear Security
PFD	probability of failure on demand
PSR	periodic safety review
SAPs	safety assessment principles
TAGs	Technical assessment guides
UKSO	UK Safeguards Office
WENRA	Western European Nuclear Regulators' Association

8 Bibliography

1 Health and Safety Executive, Nuclear Safety Assessment principles, www.hse.gov.uk/nuclear/SAPs/SAPs2006.pdf

2 Guerra, S, Bishop, P, Bloomfield, R and Sheridan, D.'Assessment and Qualification of Smart Sensors', 7th American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies NPIC&HMIT 2010, Las Vegas, Nevada, November 7-11, 2010

3 Bloomfield, R and Bishop, P 'Safety and Assurance Cases: Past, Present and Possible Future', Safety Critical Systems Symposium, Bristol, UK, 9-11 Feb 2010

4 Nuclear Installations Act 1965, www.legislation.gov.uk/ukpga/1965/57

5 Health and Safety Executive, Health and Safety at Work etc Act 1974, www.hse.gov.uk/legislation/hswa.htm

6 Health and Safety Executive, Technical Assessment Guide: Guidance on the Purpose, Scope and content of Nuclear Safety Cases, T/AST/051, Issue 001, 13 May, 2005

7 Health and Safety Executive, The licensing of nuclear installations. www.hse.gov.uk/nuclear/notesforapplicants.pdf

8 Health and Safety Executive, *License Condition Handbook*, www.hse.gov.uk/nuclear/silicon.pdf, October 2010

9 Wikipedia, Nuclear Power, http://en.wikipedia.org/wiki/Nuclear_power

10 BBC News, Windscale fallout underestimated, http://news.bbc.co.uk/1/hi/sci/tech/7030536.stm

11 Records of the UK Atomic Energy Authority and its predecessors, National Archives, www.nationalarchives.gov.uk/catalogue/displaycataloguedetails.asp?CATID=2&CATLN=1&accessmethod=5&j=1

12 Rogovin, M. *Three Mile Island: A Report To The Commissioners And To The Public.* Vol.I Washington, DC: U. S. Government Printing Office, 1980

13 Sir Frank Layfield, 'Sizewell B Public Enquiry Report', London HMSO, 1987.

14 British Energy, 'Different types of nuclear power', www.british-energy.com/documents/Different_types_of_nuclear_power.pdf

15 United Nations Scientific Committee on the Effects of Atomic Radiation, 'Consequences of the Chernobyl disaster', www-ns.iaea.org/appraisals/chernobyl.asp?s=7&l=58

16 World Nuclear Association, Chernobyl Accident, www.world-nuclear.org/info/chernobyl/inf07.html

17 'Hinkley Point B Nuclear Power Station Safety Case', www.hse.gov.uk/foi/releases/hinkleyb2.htm

18 Health and Safety Executive, 'HSE Enforcement Policy Statement', www.hse.gov.uk/foi/internalops/fod/oc/100-199/130-6.htm

19 Health and Safety Executive, 'ALARP at a glance', www.hse.gov.uk/risk/theory/alarpglance.htm

20 The Telegraph, 'Huge expansion of nuclear power signalled by Government', www.telegraph.co.uk/earth/energy/nuclearpower/6532068/Huge-expansion-of-nuclear-power-signalled-by-Government.html

21 Health and Safety Executive, 'New Reactors', www.hse.gov.uk/newreactors/background.htm

22 Littlewood, B 'The Use of Computers in Safety-Critical Applications, Final Report of the Study Group on the Safety of Operational Computer Systems constituted by the Advisory Committee on the Safety of Nuclear Installations', HSE Books London 1998.

23 IEC 61226 - 'Nuclear power plants – Instrumentation and control systems important for safety – Classification', Edition 1, 1993.

24 IEC 61226 - 'Nuclear power plants – Instrumentation and control systems important for safety – Classification', Edition 2, 2005

25 Hughes, G and Hall, RS. 'Recent developments in protection and safety-related systems for Nuclear Electric's (UK) power plant', in *IAEA/OECD International Symposium on Nuclear Power Plant Instrumentation and Control,* (Tokyo), 1992.

26 Hunns, DM and Wainwright, N. 'Software-based protection for Sizewell B: the regulator's perspective', *Nuclear Engineering International*, September, pp.38-40, 1991.

27 Pavey, D. Programmable Protection in UK NPP: 10 years on, Implementing and Licensing Digital I&C Systems and Equipment in Nuclear Power Plants' IAEA 22 to 24 November 2005, Espoo, Finland

28 Bishop PG, Chozos N, Tourlas K, 'Using Non Safety-Assured Programmable Components in Modest Integrity Systems', SAFECOMP 2010, 14-17 September, 2010, ISSN 0302-9743, (Erwin Schoitsch, Eds.), pp. pp 275-388, Springer Verlag, Vienna, Austria, 2010

29 Bishop PG, Cyra L, 'Overcoming Non-determinism in Testing Smart Devices: A Case Study', SAFECOMP 2010, 14-17 September, 2010, ISSN 0302-9743, (Erwin Schoitsch, Eds.), pp. 237-250, Springer Verlag, Vienna, Austria, 2010.

30 Lord Cullen, The public inquiry into the Piper Alpha disaster. HMSO Cm 1310, 1990

31 Lord Robens, Safety and Health at Work. Report of the Committee 1970 – 72. HMSO Cmnd 5034, 1972.

32 BS IEC 61513 ed. 1 Nuclear power plants – Instrumentation and control for systems important to safety – General requirements for systems, March 2001.

33 BS EN 61508 ed. 1 Functional safety of electrical/electronic/programmable electronic safety-related systems, 2002.

34 IAEA NS-R-1 Safety of nuclear power plants: design: safety requirements, September 2000.

35 BS IEC 60880 ed. 2 Nuclear power plants – Instrumentation and control systems important to safety – Software aspects for computer-based systems performing category A functions, May 2006.

36 International Atomic Energy Agency, www.iaea.org/

37 Ham A and Hall R. A Way Forward for Nuclear Power – 2006 Energy Review Submission, Department of Business Innovation and Skills, January 2006, www.berr.gov.uk/files/file28276.pdf

38 British Energy Generation Ltd, Sizewell B Nuclear Power Station report - Q1 2010, www.hse.gov.uk/nuclear/llc/2010/sizewellb1.htm#inspect

39 Health and Safety Executive, 'The Tolerability of Risk from Nuclear Power Stations', 1992, www.hse.gov.uk/nuclear/tolerability.pdf

Appendix D1

In Section 2.1.4.3 we identified a set of key international standards. The following extracts are key points from some standards and guidelines and reflect UK practice regarding claim limits for software. IAEA refers to an early version of IEC 61226 (from 1993),[23] but the software claim limit is reiterated in the revised version of IEC 61226 (2005).[24]

IAEA NS-G-1.1 (Safety and Reliability Issues, Section 2.9)

Hence, designs requiring a single computer based system to achieve probabilities of failure on demand of lower than 10–4 for the software should be treated with caution (cf. para. 8.2.2 of).

IEC 61226 – 2005 (Section 7.3.2.1)

For an individual system which is specified and designed in accordance with the highest quality criteria, a figure of the order of 10–4 failure/ demand may be an appropriate overall limit to place on the reliability that may be claimed, when all of the potential sources of failure due to the specification, design, manufacture, installation, operating environment, and maintenance practices, are taken into account. This figure includes the risk of common mode failure in the redundant channels of the system, and applies to the whole of the system, from sensors through processing to the outputs to the actuated equipment. Claims for better reliabilities than this are not precluded, but will need special justification, taking into account all of the factors mentioned. Alternatively, the design of independent I&C systems important to safety with an acceptable level of diversity may be applied.

TAG 46 (Appendix 3)

Taking all of the above into account, a probability of 1E-4 failures per demand is considered to be the best that can justifiably be claimed for computer based software systems used in circumstances where the consequence in the event of failure of the software system could potentially involve very large releases of radioactive material. However, it is recognised that advances in system design and software engineering techniques might lead to a situation where a strong case could be made for a lower figure. Such a case would not then be ruled out of consideration.

Supplement E:

Safety case use in the petrochemical industry

Jamie Henderson
Human Reliability Associates

This is one of a series of supplements to the report: *Using safety cases in industry and healthcare: a pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare*

To access the report and the other supplements, please visit www.health.org.uk/safetycasesreport

Contents

Supplement E:

Safety case use in the petrochemical industry

1 Introduction

This summary is written from the perspective of the European, and specifically the UK, petrochemical industry, as these are the territories with which the authors are most familiar.

In the petrochemical industry, safety cases are primarily developed to assist in the prevention of major accidents. A major accident is defined as:

> *'…an occurrence (including in particular, a major emission, fire or explosion) resulting from uncontrolled developments in the course of the operation of any establishment and leading to serious danger to human health or the environment, immediate or delayed, inside or outside the establishment, and involving one or more dangerous substances' – Control of Major Accident Hazards (COMAH) Regulations (1999)*[1]

These events have a low probability of occurrence but their consequences are potentially high, with the possibility of multiple fatalities both on and off site. Historically, analyses in safety cases for installations have focused on these types of events. One review of high-cost petrochemical accidents, covering the 22 years since the 1974 Flixborough disaster, identified 20 accidents with a cost of £1.5 million or greater (at 1996 prices) in the UK alone.[2]

However, until relatively recently, safety metrics in the UK petrochemical industry have been targeted almost exclusively at what has become known as **occupational safety** (eg injuries arising from slips, trips and falls, or manual handling).

One possible reason for this measurement focus is that these types of event have a relatively high probability of occurrence, making it easier for sites to meaningfully measure performance. For example, sites often have signs indicating how many days have passed since a lost time injury. A realisation that such measures may not provide the best indication of whether a site is likely to suffer a major accident has resulted in a new effort to identify and measure precursor indicators of major accidents.[3]

2 Regulatory context and best practice

In the UK, the legislation that governs the requirements related to safety cases in the petrochemical industry is different for onshore and offshore facilities. Offshore installations are subject to the Safety Case Regulations (SCR),[4,5] onshore establishments are covered by the Control of Major Accident Hazards Regulations (COMAH). The SCR came into force following the Lord Cullen inquiry into the 1988 Piper Alpha disaster.[6] There are many similarities between the SCR and COMAH legislation (unfortunately, however, the terminology is different – the SCR legislation refers to safety cases, whereas the COMAH legislation refers to safety reports). Therefore, in the interests of brevity, this summary focuses on the COMAH legislation.

The Control of Major Accident Hazards Regulations 1999 came into force on 1 April 1999 and were amended by the Control of Major Accident Hazards (Amendment) Regulations 2005.[7] COMAH regulations implement Council Directive 96/82/EC, known as the Seveso II Directive, as amended by Directive 2003/105/EC and replaced the Control of Industrial Major Accident Hazards Regulations 1984 (CIMAH). This section sets out the requirements for safety cases in the context of this legislation. Issues related to previous legislation and the historical development of this legislation in a European context are discussed in the following section. The information in this section is, unless otherwise stated, a summary of information taken from the legislation.

In the UK, the competent authority (CA) is responsible for the enforcement of the COMAH regulations. The CA comprises three organisations: the Health and Safety Executive (HSE), the Environment Agency (EA) and the Scottish Environmental Protection Agency (SEPA). The HSE website[8] describes how the CA works in the context of UK health and safety legislation, which is based on the principle that 'those who create risk are best placed to manage it', and sets out responsibilities for three key stakeholders:

- **Operators of COMAH sites** – have a responsibility to manage risks arising from their activities.
- **Emergency planning authorities** – have a duty to ensure that off-site emergency plans are prepared and adequate.
- **Competent authority** – assesses safety reports submitted by site operators, inspects sites and investigates incidents. It has the power to prohibit operation where there is evidence that measures taken for prevention and mitigation of major incidents are seriously deficient. The CA charges the operators fees for services.

In the COMAH regulations, sites are designated as top-tier or lower-tier based on the inventories of dangerous substances they hold. The threshold quantities for each tier are specified in the regulations. Therefore, for example, a site with an inventory of more than 500 tonnes of methanol would be a lower-tier site, whereas a site holding more than 5,000 tonnes would be top-tier. Currently, there are around 650 lower-tier sites and 375 upper-tier sites in the UK.[9]

In COMAH, safety cases are referred to as safety reports. Only top-tier sites are required to complete a full safety report. However, lower-tier sites are required to develop a major accident prevention policy, including an overview of their safety management system. The authors are currently working with one lower-tier site that has developed a full safety report despite their inventories not requiring it. They have chosen to do this as their operating context sees them interacting with top-tier sites. The main purposes of a safety report, as set out in the legislation, are to demonstrate that:

- an installation-specific major accident prevention policy, and a safety management system for implementing it, are in place
- major accident hazards have been identified and that the necessary measures have been taken to prevent such accidents and to limit their consequences
- adequate safety and reliability principles and safeguards have been incorporated into the design, construction, operation and maintenance of an installation
- on-site emergency plans have been developed, and that information has been supplied to enable the development of off-site emergency plans.

The principal requirements for COMAH sites, in addition to a general obligation to take all measures necessary to prevent major accidents and limit their consequences, are summarised in Table E1, overleaf.

Table E1: Principal requirements for COMAH sites

Requirement	Notes
Produce a Major Accident Prevention Policy (MAPP)	The MAPP should: – be proportionate to the hazards presented by the establishment. – set out aims and principles of action with respect to major accident hazards (MAHs). – demonstrate that a safety management system (SMS) is in place to determine and implement the MAPP (including coverage of the following: roles and responsibilities, identification and evaluation of MAHs, operational control, management of change, planning for emergencies, monitoring performance and audit and review).
Notify the competent authority	Sites are required to specify the dangerous substances present at the site and describe the activities that are planned to take place.
Produce a safety report*	For top-tier sites, the safety report should contain the MAPP, including the safety management system demonstration, as well as the following information: – A description of the installation (including, for example, geographical location, activities which could result in a major accident, a description of the processes, an inventory of dangerous substances). – A detailed description of possible major accident scenarios (including their probability, a summary of triggering events, an assessment of severity and consequences, a description of technical parameters and safety equipment). – A description of protection and intervention measures (including equipment, alerting systems, and mobilisable resources). – Safety reports should be reviewed at least every five years, or when new facts emerge, or whenever the operator makes a change to the safety management system.
Produce an emergency plan*	Top-tier sites should produce an on-site emergency plan that sets out: – roles and responsibilities – for foreseeable events, a description of actions which should be taken to control and limit their consequences – how warning systems are expected to function – the interface with the local authority responsible for setting the off-site plan in action – arrangements for training staff in emergency duties. In addition, an off-site emergency plan should also be developed by the local authority. The legislation requires that both parts of this plan should be reviewed every three years and that the plan should be tested.
Provide information*	Top-tier sites are required to: – provide information to the public (eg regarding how the population will be warned of a major accident) – provide information to the competent authority (eg in response to requests or to inform them of a major accident).

* These steps are only required for top-tier sites.

The predictive elements of a COMAH safety report should demonstrate that measures have been taken to reduce the likelihood of hazards, and to mitigate their consequences, until the associated risks are as low as reasonably practicable (ALARP). In general, the greater the hazard, and associated risks, the more the operator of a facility should favour adopting additional measures, unless a cost-benefit analysis suggests otherwise. The HSE's Safety Report Assessment Manual[10] states:

> *'If all reasonably practicable measures are in place, and the risks are tolerable, then there is nothing more to be done – Individual Risk and Societal Concern must be ALARP.'*

The requirement for the safety report is that it should provide evidence and a depth of argument that is proportionate to the risk posed by the establishment. More detail regarding ALARP and tolerability of risk is provided in the HSE guidance document, *Reducing Risks, Protecting People*.[11]

In terms of measures specified in safety reports for managing risk, a key standard is IEC 61511[12] (This is the process industry-specific version of IEC 61508:1998[13]). This is concerned with the functional safety of safety instrumented systems (SIS) that might, for example, be specified as a risk control measure to automatically shut down a process operation, in the event of an abnormal situation, to prevent a hazard being realised. The HSE *Safety Report Assessment Manual* (SRAM) requires that a selection of SIS linked to major accident scenarios should be provided in the safety report.

3 Development and drivers

In Europe, the history of safety cases and safety case legislation has been closely linked to the history of high-profile accidents.

In particular, the 1976 Seveso accident, at a small chemical manufacturing facility in Italy, resulted in the development of legislation aimed at the prevention and control of major accidents. A simple chronology is presented in table E2 below.[14]

Table E2: Chronology of accidents

Date	Event	Notes and relationship to safety case requirements
1974	Flixborough disaster[15]	Large explosion kills 28 workers. Failure to use systematic analysis process to identify hazards (eg HAZOP) identified by the HSE investigation.[16]
1976	Seveso accident[16]	An uncontrolled exothermic reaction results in the release of a dense vapour cloud containing poisonous and carcinogenic dioxin. Ten square miles of land are contaminated, more than 600 people are evacuated and 2,000 treated for poisoning.
1982	Seveso Directive is adopted	Council Directive 82/501/EEC on the major accident hazards of certain industrial activities – the so-called Seveso Directive – is adopted.[17] Required substances to be identified and processes described. No requirement to include MAPP or safety management system.
1984	Bhopal disaster[18]	A leak of gas and other chemicals from a plant in India resulted in the exposure of hundreds of thousands of people. Estimates on the death toll varied from 2,000 to as many as 15,000 people.
1984	Control of Industrial Major Accident Hazards (CIMAH) regulations adopted in UK for onshore facilities	Superseded by COMAH regulations in 1999. Similar to Seveso I requirements with an emphasis on description.

Date	Event	Notes and relationship to safety case requirements
1987/ 1988	Two revisions of the Seveso Directive	These amendments were aimed at broadening the scope of the Directive to include the storage of dangerous substances.
1988	Piper Alpha disaster[6]	An oil platform that was later converted to gas production. An explosion on the platform and the resulting fire killed 167 men with only 59 survivors.
1992	Safety case regulations adopted for UK offshore industry	The publication of Lord Cullen's report into the Piper Alpha disaster in 1990 paved the way for the introduction of formal safety case requirements in the UK offshore industry.
1996	Seveso II Directive is adopted	Implemented in the UK as the COMAH regulations (see below).
1999	Control of Major Accident Hazards (CIMAH) regulations adopted in UK for onshore facilities	Replaced the CIMAH regulations and introduced a greater degree of uniformity with the offshore SCR. The regulations brought a number of smaller sites under the legislation and introduced a number of new features, including the MAPP and safety management system requirements. Also brought an increased emphasis on demonstration rather than description.
2003	Revision of Seveso II Directive	Revision of Seveso II directive to include additional requirements for risk assessment. The most important extensions of the scope cover risks arising from storage and processing activities in mining, from pyrotechnic and explosive substances, and from the storage of ammonium nitrate and ammonium nitrate-based fertilizers[14].
2005	Revision of safety case regulations for UK offshore industry and COMAH regulations	

The evolution of the safety report requirements in the UK has not been without difficulty. In particular, the change from description under the CIMAH legislation to demonstration under the COMAH legislation caused some confusion. This requirement for demonstration placed the onus on operators of facilities to show that they were managing their risk using the ALARP concept. The HSE's own website acknowledges some of the difficulties with this approach:

'Using "reasonably practicable" allows us to set goals for duty-holders, rather than being prescriptive. This flexibility is a great advantage but it has its drawbacks, too. Deciding whether a risk is ALARP can be challenging because it requires duty-holders and us to exercise judgement. In the great majority of cases, we can decide by referring to existing "good practice" that has been established by a process of discussion with stakeholders to achieve a consensus about what is ALARP. For high hazards, complex or novel situations, we build on good practice, using more formal decision making techniques, including cost-benefit analysis, to inform our judgement.'[19]

Initially, this was addressed in different ways by operators (eg in terms of the maximum level of risk, in terms of cost-benefit analysis, or in terms of the technology adopted) and consequently, the guidance on this concept, provided by the competent authority, evolved after the submission of the first batch of COMAH safety reports.[20] This issue illustrates the types of challenges that can be faced by the authorities made responsible for assessing and responding to safety cases.

4 Types of safety cases and content

To summarise, a typical COMAH safety report will contain the following elements:

- major accident prevention policy (MAPP) (including safety management system description)
- an identification of hazards
- analysis and assessment of risk (including description of prevention/limitation measures)
- an emergency plan.

The COMAH regulations are goal-setting rather than prescriptive, in terms of the methods that should be used to make these demonstrations. Therefore, a given safety report may include a range of different methods. The following sections provide some examples of typical methods and approaches that are used for the predictive aspects of the report (ie the identification of hazards and the risk assessment).

Identification of hazards

The COMAH *Safety Report Assessment Manual* (SRAM)[5] provides some indication of the types of approach that might be seen as acceptable for this purpose:

- Hazard and operability studies (HAZOP).
- Safety review and studies of the causes of past major accidents and incidents.
- Industry standard or bespoke checklists for hazard identification.
- Failure mode and effect analysis (FMEA).
- Job safety analysis (eg task analysis).
- Human error identification methods.

Whatever approach is used, firstly, according to the SRAM, the report should demonstrate a clear understanding of the site operations, the materials involved and the process conditions. Secondly, it must identify the hazards to people on-site and off-site and to the environment. Finally, there must be an analysis of different ways the hazard can be eliminated, reduced in scale, realised and controlled.

There is a requirement that all foreseeable major accidents are identified by the hazard identification exercise. However, these may be organised into groups to make the risk assessment feasible. If this is the case, the identified scenarios should be shown to dominate the risk and should include worst-case scenarios.

Analysis and assessment of risk

Following the hazard identification the SRAM suggests that, for the hazards that remain, the following activities should take place.

- There must be a prediction of the likelihood of the hazards being realised, taking account of the chance of success and failure of possible preventative measures.
- A prediction of the corresponding consequences both when the mitigation measures work and fail.
- An analysis of the risks associated with the remaining hazards and the options for reducing them.
- A decision about which measures need to be implemented to make the risks to people (individually and collectively) and the environment satisfy the ALARP criteria.
- A presentation of the results of the risk assessment to provide the evidence and arguments which demonstrate that all measures necessary have been taken to prevent and mitigate major hazards.

The COMAH regulations do not state whether qualitative, semi-qualitative or quantitative arguments should be used. Instead, there is a requirement for the operator to use methods that are most appropriate for the risk involved. The SRAM suggests some techniques that may be used to develop frequency and probability estimates:

– Relevant operational and historical data.

– Fault tree analysis (FTA).

– Event tree analysis (ETA).

IEC 61511[10] provides useful detail for managing risk in relation to the concept of protection layers. Part 3: Annex C of this guidance describes how protective systems can be characterised as **safety layers**. In this characterisation, the process has a number of protection layers (ie the process control system, alarms, SIS, relief devices, physical protection and emergency response) that combine equipment and administrative controls to control and mitigate process risk. Annex F describes layer of protection analysis (LOPA), a semi-qualitative method, as a formal process for considering initiating causes (eg as identified in a HAZOP study) and the protection layers that prevent or mitigate the hazard. This process can be used to determine the total existing amount of risk reduction and the need for further risk reduction. The LOPA approach allows for the identification of the appropriate safety integrity level (SIL) for the safety-instrumented function.

5 Discussion

It seems that there are two issues of concern: the first is whether or not safety cases have made an improvement to safety performance in petrochemical industries. If this is the case, then the next question is the degree to which the type of safety case regime in place for the petrochemical industries would be transferable to a healthcare setting. This issue is addressed in the next section.

With regard to the first question, there is little hard evidence to support the proposition that safety cases have made an improvement to safety performance in the petrochemical industry. We are unaware of any research that demonstrates such a link. However, their application, in much the same way as for accident investigation, has high face validity. Certainly, it would be difficult to argue that reviewing hazards and assessing risks could have anything other than a positive impact on safety performance.

However, safety report development is a potentially onerous activity for the establishments concerned. This is particularly the case in a goal-setting environment where the onus is on the operators themselves to demonstrate that they are managing risk. There may be a danger, for example, of an organisation employing an external consultancy to develop a standard report that does little to engage the operator in the review of risks at their facility.

A continuing challenge is dealing with the low probability nature of the major accidents that the safety report seeks to address. The industry is currently engaged in an effort to identify process safety indicators that may more accurately indicate sites that are vulnerable to major accidents.[3]

6 Lessons and recommendations for healthcare

The predominant historical driver for the initiation and development of a safety case regime in the petrochemical industry has been major accidents (eg Seveso, Bhopal, Flixborough, Piper Alpha). These accidents have often had an impact beyond the immediate facility and been extensively covered by media outlets. This, in turn, has prompted a wider discussion about the risks society is willing to tolerate in exchange for the benefits that such endeavours bring. In the UK, for example, the HSE has published its criteria for decisions related to the regulation of risk in the context of changes in the preferences, values and expectations of society.[9] The healthcare sector, however, does not have the dubious benefit of such high-profile incidents to drive the requirement for safety cases. Major accidents in the petrochemical industry are, for want of a better description, often visually spectacular events with the potential for widespread harm and damage to the environment. Major accidents in a healthcare setting are most likely to harm an individual and take place, potentially, in a context of ill-health where deterioration is not unexpected. This, of course, may change, and there are examples of high-profile medical accidents. However, if high-profile accidents really have been the primary driver for the adoption of these techniques in the petrochemical industry, then healthcare may have to find other ways to promote a requirement for safety cases.

Safety cases are a regulatory requirement, but they also help organisations develop confidence that they are adequately managing risks at their sites. In a healthcare setting it may be unlikely, at least in the short term, that there will be a similar regulatory push. Therefore, the benefits

to healthcare organisations would need to be obvious in order to drive such a process. This may be difficult; our anecdotal experience in the petrochemical industry is that regulation is the main influence on take-up of safety case analyses. This is not to say that there is a lack of interest in safety performance; the vast majority of individuals have a keen, motivated interest in the topic. However, these are complicated businesses that have numerous goals related to production and safety. Therefore, the activities that are pursued tend to be those which have the highest priority and, unsurprisingly, regulatory pressure is a significant factor in determining priority. This is particularly true of more novel aspects. For example, human factors issues are now addressed by specialist inspectors from the competent authority. Until this happened, our experience was that site managers would appreciate the issues related to human factors, but would be unwilling to divert resources to their management, as they had numerous other priorities. This altered as soon as the regulator became clear that they were serious about requesting a systematic analysis of human factors issues related to major accident hazard-critical tasks.

A further significant challenge is posed by the resources required to develop COMAH safety reports. Ideally, these reports would be living documents; continuously maintained as a risk register and a demonstration that the site is managing safety. However, the requirement for demonstration in relation to the ALARP concept means that reports often run to several hundred pages. They are typically complex documents that take up considerable site time and effort to develop. Moreover, many sites use external consultants to assist with the report development. It is unlikely that such a level of effort would be possible to sustain without external regulatory pressure. There is also the question of whether such a focus on the event of safety report development is even desirable. For example, it is not difficult to imagine a situation where the development of the safety report detracts from the day-to-day running of the plant.

Under the COMAH regulations, the requirement to complete a safety case is determined by an analysis of the hazards present at a given site. In the petrochemical industry this is relatively straightforward, as an assessment of inventories of hazardous substances can readily be performed. However, for healthcare, such a simple determination of hazards is less straightforward. Without a detailed knowledge of the healthcare environment, one supposes that harm is most likely to be realised during procedures performed on patients. This would suggest that certain procedures with high potential harm might be classified in a similar manner. However, if most procedures are generic to all hospitals, then such a classification may not add value.

In addition, if hospitals, for example, were asked to identify their most hazardous procedures, then these could potentially be pre-identified (eg in the same way that named substances are listed in the COMAH regulations). Hospitals could then be required to demonstrate management of risks in relation to such procedures.

Regulation in the UK petrochemical industry is goal-setting rather than prescriptive. This means that sites are free to choose the analysis methods that they feel are most appropriate to demonstrate their management of risk. One downside of this is that more effort is required – both from the operator, in terms of the safety case development, and the regulator, in terms of the review of the report and the inspection of the facilities.

7 Glossary

Term		Notes
ALARP[21]	As low as reasonably practicable	A concept involving weighing a risk against the trouble, time and money needed to control it.
FMEA[22]	Failure modes and effects analysis	A technique for evaluating the ways in which equipment can fail (or be improperly operated) and the effects these failures can have on a process.
HAZOP[23]	Hazard and operability study	A systematic method for the identification of hazards and problems that might prevent efficient operation of a plant.
LOPA[24]	Layers of protection analysis	A technique that can be used to assess the adequacy of the layers of protection provided for an activity. This includes a risk analysis that can be used to identify safety functions that need to be put in place to reduce risk to a tolerable level.
SRAM[10]	*Safety Report Assessment Manual*	Manual used by HSE inspectors when assessing COMAH sites.

8 References

1 The Control of Major Accident Hazards Regulations 1999 No 743. Overview available at: www.hse.gov.uk/comah/background/comah99.htm – accessed February 2011.

2 Fewtrell, P. and Hirst, I.L. (1998) 'A review of high cost-cost chemical/petrochemical accidents since Flixborough 1974.' *IChemE Loss Prevention Bulletin*. No 140.

3 Health and Safety Executive (2006) *HSG254 Developing process safety indicators: A step-by-step guide for chemical and major hazard industries*. HSE Books.

4 The Offshore Installations (Safety Case) Regulations 1992 No 2885. Available at: /www.legislation.gov.uk/uksi/1992/2885/contents/made – accessed February 2011.

5 The Offshore Installations (Safety Case) Regulations 2005 No 3117. Available at: www.legislation.gov.uk/uksi/2005/3117/contents/made – accessed February 2011.

6 Cullen, The Honourable Lord (1990) *The Public Inquiry into the Piper Alpha Disaster*. HMSO.

7 The Control of Major Accident Hazards (Amendment) Regulations 2005 No. 1088. Available at: www.legislation.gov.uk/uksi/2005/1088/contents/made – accessed February 2011.

8 Health & Safety Executive. *The Competent Authority in Detail.* Available at: www.hse.gov.uk/comah/authority – accessed February 2011.

9 Personal communication with HSE Hazardous Installations Directorate, February 2011.

10 Health and Safety Executive (2006) *COMAH Safety Report Assessment Manual (V2).* Available at: http://www.hse.gov.uk/comah/sram/ – accessed February 2011.

11 Health and Safety Executive (2001) *Reducing Risks, Protecting People: HSE's decision-making process*. HSE books.

12 BS IEC 61511:2003 Functional safety – Safety Instrumented Systems for the process industry sector.

13 BS IEC 61508 1998 Functional Safety of electrical/electronic/programmable electronic safety-related systems.

14 Chemical Accidents (Seveso II) – Prevention, Preparedness and Response. European Commission. Available at: http://ec.europa.eu/environment/seveso/index.htm – accessed February 2011

15 Report of the Court of Inquiry (1975) *The Flixborough Disaster*. HMSO 1975

16 Health & Safety Executive. Flixborough (Nypro UK) Explosion 1st June 1974. Available at: www.hse.gov.uk/comah/sragtech/caseflixboroug74.htm – Accessed February 2011.

17 European Commission: Environment. *Chemical Accidents (Seveso II) - Prevention, Preparedness and Response.* Available at: http://ec.europa.eu/environment/seveso/index.htm – accessed February 2011.

18 Morehouse, W. and Arun, M. (1986) *The Bhopal Tragedy*. The Council on International and Public Affairs: New York.

19 Health and Safety Executive. *ALARP at a glance*. Available at: www.hse.gov.uk/risk/theory/alarpglance.htm – accessed February 2011.

20 Ennis, T. (2003) *CIMAH to COMAH: Demonstration to implementation*. Presentation made at the 34th meeting of United Kingdom Explosion Liaison Group. Available at: http://ukelg.ps.ic.ac.uk/34TE.pdf - accessed February 2011

21 Health and Safety Executive. *ALARP at a glance*. Available at: www.hse.gov.uk/risk/theory/alarpglance.htm - accessed February 2011.

22 Center for Chemical Process Safety (1992) *Hazard Evaluation Procedures*. American Institute of Chemical Engineers: New York.

23 Kletz, T. (1999) *HAZOP and HAZAN*. Institution of Chemical Engineers, Rugby.

24 Center for Chemical Process Safety (2001) *Layer of Protection Analysis – Simplified Process Risk Assessment*, American Institute of Chemical Engineers: New York.

Supplement F:

Safety case use in the railway industry

John Medhurst and David Embrey
Human Reliability Associates

This is one of a series of supplements to the report: *Using safety cases in industry and healthcare: a pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare*

To access the report and the other supplements, please visit www.health.org.uk/safetycasesreport

Contents

Supplement F:

Safety case use in the railway industry

1 Introduction

This review is written primarily from the perspective of the UK railway industry, although recent European legislation is also considered.

Safety cases in most high-risk industries are developed to assist in the prevention of major accidents – those involving fatalities, whether to staff, passengers or both. Railways are an extremely safe mode of transport, and in the past 50 years, multiple fatality-level accidents involving passengers in the UK have occurred at relatively infrequent intervals compared with other transport modes.[1,9]

The major influence toward the adoption of a safety case regime for regulation of the railway industry in the UK was the run-up to privatisation of the railways in 1994.[2] The UK managed railway safety under a safety case regime from 1994 until 2006, at which point a certification-based system was introduced to comply with a European-wide legislative initiative,[3] which focused on interoperability.

The railway regulatory regime in the UK in the past 20 years has focused on both occupational and passenger (general public) safety. It has also been closely involved with construction industry safety, since large rail infrastructure projects are major projects in their own right (for example, the London Crossrail project is said to be the largest current infrastructure project of any kind in Europe).

2 The regulatory framework

This section considers the regulatory framework for railways that was in place in the UK between 1994 and 2006, which included safety case legislation. There was a single regime for all major railways (excluding only some light rail and 'heritage' systems), with differences in application of the legislation for infrastructure controllers and others (typically train and station operators). The original applicable regulations were the Railways (Safety Case) Regulations 1994, which were updated as the Railway (Safety Case) Regulations 2000, with further amendments in 2001 introduced by the Railway Safety (Miscellaneous Amendments) Regulations 2001. These will hereafter be described collectively as the R(SC)Regs or the Regulations. These regulations were all made under the enabling legislation of the Health and Safety at Work Act 1974.

In the UK, the competent authority (CA) responsible for the enforcement of the R(SC)Regs was Her Majesty's Railway Inspectorate (HMRI), previously a separate agency but which became part of the Health and Safety Executive in 1993. HMRI was transferred to the Office of the Rail Regulator (ORR) in April 2006, concurrently with the new legislation implementing the European Railway Safety Directive.[4]

To operate as a railway operator of any description – ie an infrastructure controller, train or station operator – a company was required to obtain the CA's formal approval of a safety case. Infrastructure controllers (ICs) were required to obtain a report from an assessment body recommending their safety case for approval before it was sent to the CA for approval.

A train or station operator who was not also the IC, was required to send its safety case to the IC and also to an assessment body before forwarding it to the CA for review. This two-tier review process was the feature that distinguished the regulatory framework for infrastructure controllers from that of operators. The concept behind this was to ensure an integrated system-wide review of safety for rail operations managed jointly by different organisations.

Until late in 2001, the single infrastructure controller for the whole surface rail system of any significance was Railtrack, a company set up as a plc at privatisation. Railtrack went into liquidation in November 2001 and was taken into private ownership, emerging as a restructured not-for-profit private company limited by guarantee in October 2002. The stakeholders of Network Rail are mainly rail operating and maintenance companies, but the Department for Transport is also a stakeholder and has special membership rights. The operators are all private companies (with the exception of some publicly owned urban railways such as London Underground running trains on surface rail infrastructure) and, although substantial, are very much smaller than Network Rail.

The principal duties and functions under the Regulations were as shown in Table F1 below.

Table F1: Functions and duties under the UK Railway safety case regime (2000 Regulations)

Functions	Duties
Train and station operators	Submit safety cases and material revisions to CA and the IC. Undertake a thorough review of the safety case at least every three years (or as directed by the CA).
Infrastructure controllers (ICs)	Submit safety cases and proposed revisions to CA and the IC. Obtain independent assessment for review by CA. Undertake a thorough review of the safety case at least every three years (or as directed by the CA) and submit proposed revisions. Assess train and station operators' safety cases and obtain assessment body review, then send own and assessment body's assessments and recommendations to CA. Commission an assessment body to undertake annual audits of their own and operators' health and safety management Systems. Ensure compliance and notify CA of any failures to comply with IC or operators' safety cases.
Assessment body	Assess IC and operators' safety cases and undertake annual audits of their health and safety management systems and make recommendations for improvement where appropriate.
Competent authority (HSE/HMRI)	Formally accept (or reject with reasons) all safety cases under the Regulations. Direct ICs or operators to revise their safety cases. Direct an IC to make a recommendation based on an operator's SC. May grant individual or generic exemptions, with conditions as appropriate.
Secretary of State (Department of Transport)	Determine appeals from decisions of the CA. Appoint persons to decide appeals on the Department's behalf.

As in the petrochemical industry, the purpose of the safety case is to demonstrate that all risks have been identified and managed, with suitable risk control measures in place to ensure that they are reduced to a level which is as low as reasonably practicable, the ALARP test, which is now standard across high-risk industries (see Supplement E).

Following extensive development work by the European Commission (EC), particularly DG VII (Transport), the 2004 Rail Safety Directive mandated the UK government to modify the 'permissioning' safety case system and introduce a Europe-wide certification system for assuring safety, thus ensuring interoperability of trains crossing one or more national borders within the Community.

3 Development and drivers

In Europe, the history of safety cases and safety case legislation in high-risk industries has been closely linked to the occurrence of high-profile accidents. In the UK, the main driver for the introduction of a safety case regime was privatisation, but several high-profile accidents gave rise to innovations in the management of safety by the railway industry that enabled it to introduce a safety case regime as the next logical step.

In particular, the 1987 fire on an escalator at King's Cross on the London Underground and the 1988 Clapham mainline derailment led to public inquiry reports that implemented fundamental changes in the assessment of risk and management of safety on urban and mainline railways.

Table F2: Railway accidents and safety case regime development

Date	Event	Notes and relationship to safety case requirements
1975	Moorgate underground train disaster – train overrides terminal station buffers – 43 deaths	Significant investment in improvement of protection for trains entering terminal stations, mainly at the hardware level.
1987	King's Cross station fire (London Underground) – 31 deaths	Radical reform of management on the Underground, including the introduction of safety management systems and the first system-wide quantified risk assessment (by 1991).
1988	Clapham derailment – 35 deaths	Major reforms within British Rail reflecting the response to King's Cross on London Underground.
1991	European Directive 91/440/EC[5] requires rights of access for international rail traffic	Coincides with preparations for the privatisation of railways in the UK and the separation of infrastructure and operations.
1992	UK government white paper announcing formal proposals for the privatisation of British Rail	The principal driver for the subsequent safety case regime.
1993	HSE report 'Ensuring safety on Britain's Railways' makes proposals for rail safety case regime, to meet the challenges of privatisation	This report considered various options and set out the proposals for the selected safety case regime in substantially the same form in which it was implemented in 1994.

Date	Event	Notes and relationship to safety case requirements
April 1994	Privatisation of British Rail	Concurrent with introduction of safety case regime (see below).
April 1994	Enactment of the Railways (Safety Case) Regulations, 1994.	See above. First introduction of a mandatory safety case regime in the UK.
October 1994	Cowden rail accident – 5 deaths	Too soon after the introduction of safety case regulations for them to be considered to have had any effect.
September 1997	Southall collision – 7 deaths	Signal operated by the infrastructure controller passed at danger by a driver employed by a train operating company
October 1999	Ladbroke Grove collision and fire – 31 deaths	Also a signal passed at danger (SPAD) incident. Southall and Ladbroke Grove accidents led directly to (*inter alia*) a review of the safety case regime.
2000	Enactment of Railways (Safety Case) Regulations 2000 and 2001 amendments, revising the safety case regime	New regulations directly reflect the analysis and recommendations of the inquiries into Southall and Ladbroke Grove.
October 2000	Hatfield derailment – 4 deaths	Derailment due to failure of track (permanent way). Infrastructure controller convicted of breaches of health and safety legislation.
2000 – 2001	Part-privatisation of the London Underground ('leasing' infrastructure stewardship to a public private partnership (PPP))	Not a success commercially and terminated in 2010 due to bankruptcy of Metronet (one of the PPP consortia), but no major accidents during the nine years of PPP infrastructure management.
November 2001 – October 2002	UK railway infrastructure owner Railtrack becomes insolvent. Assets transferred to a new company, Network Rail	A dramatic event that created a great deal of uncertainty over the future of the privatised railway.
May 2002	Potters Bar derailment – 7 deaths	Track maintenance contractor narrowly escaped prosecution.
2004	EC Railway Safety Directive	Move to a system of certification, driven by European requirements for interoperability.
2006	New regulatory framework in the UK including the Railways (Safety) Regulations 2006	Implementation of the 2004 EC Directive and an end to the 'permissioning' safety case regime for railways in the UK.[6]

Table F2 omits certain accidents, such as the Cannon Street derailment (1990, two deaths), where the fatalities were almost certainly caused by the poor crashworthiness of obsolete rolling stock, and the Great Heck derailment and collision in 2001, which was initiated by a road vehicle.

It will be seen that the introduction of privatisation to the surface railways in the UK and part-privatisation on the London Underground led to a period of commercial instability and insolvency in both cases. However, the introduction of the safety case regime proceeded relatively smoothly, as the industry had already implemented most of the groundwork (particularly risk assessment, safety management systems and the development of railway specific standards) in the late 1980s and early 1990s.

As in the petrochemical industry, the railway industry struggled with the concept of ALARP initially. In railways, it is assumed to have a techno-economic justification, with the concept of the 'cost of avoiding a fatality' playing a key part of the cost-benefit calculation. However, the industry has developed clear standards setting out procedures for carrying out ALARP calculations. These were largely developments of already existing public sector accounting methods, which allowed social and safety benefits to be taken into consideration, including the cost of avoiding a fatality.

4 Types of safety cases and content

There were two distinct classes of safety case under the UK railway regime, those for infrastructure controllers and those for other operators. The legislation imposed on ICs a duty to review the safety cases of operators using their infrastructure, in conjunction with the IC's own safety case.

As with COMAH, the Railway (Safety Case) Regulations were goal-setting (or 'permissioning' as the HSE described it latterly) rather than prescriptive, in terms of the methods that should be used to make the required demonstrations.

The key elements of a railway safety case were as follows.

1. Details of the safety case 'duty holder' and a description of its operation.
2. Details of risk assessments, including methodology, results and the implementation of risk control measures.
3. Description of the health and safety management system and a demonstration of its effectiveness, including provisions for audit and review.
4. Description of the technical specifications for plant and equipment.
5. Description of the operational and maintenance procedures.
6. Particulars of arrangements for records of training and competence of staff.
7. Arrangements for cooperating with other railway bodies.
8. Arrangements for incident investigation and emergency response, including evacuation of stations and trains.
9. Safety case development plan.

Some key elements of the safety case regime are discussed further in the following sub-sections.

Identification of hazards

As in other high-risk industries, a standard family of techniques is used to identify hazards. Hazard and operability studies (HAZOP), failure modes and effects analysis (FMEA) and various forms of human reliability assessment are commonly used. The HAZOP technique is used particularly to deal with the interface between technology and operations and has the advantage that the HAZOP workshops bring together technicians and operations staff who might not otherwise interact to consider problems in the same depth.

FMEA is used at the hardware and software level to predict potential failures, with the particular goal of designing out 'single-point' failures. FMEA is routinely extended to full reliability, availability, maintainability and safety (RAMS) studies following recognised international standards, which effectively develop a generic safety case for a particular item of equipment in a system context.[7] Most mechanical/electrical/electronic systems lend themselves to this approach which is now in wide use throughout industry, not just in railways.

Particular emphasis is placed on consideration of operations in normal, degraded, abnormal and emergency conditions, to ensure that all hazards are identified, particularly in an operational context. Human factors error analysis techniques are particularly useful in assessing such scenarios – for example, train and station evacuation.

Analysis and assessment of risk

Risk assessment typically follows the traditional engineering industry process of using fault tree analysis (FTA) to determine the failure rate or conditional probability of a mechanical or system failure, followed by event tree analysis (ETA) to calculate the consequences of failure in terms of fatalities and major / minor injuries. Human error assessment techniques are now routinely used to assess probabilities of base events in the fault trees or nodes of the event trees. However, the guidance notes to the R(SC) Regs[8] emphasised that quantified risk assessment (QRA) should be treated as an aid to decision-making only and not a complete substitute for qualitative methods.

Statistics are maintained of base event failure rates. Where no figures exist, these are assessed using formal techniques such as FMEA or its cousin FMECA (the C is for criticality) or are assessed by domain experts.

The effectiveness of risk control measures is tracked by calculating the effect of implementation on failure and fatality rates.

Safety management

Great emphasis is placed on safety management systems (SMS) in the management of railway safety. The SMS was the first part of the modern safety regime to be put in place in the late 1980s and there is still a tendency to conduct the whole of the management of a safety case regime from the structure of the SMS. Thus the SMS often encompasses activities such as failure analysis, incident investigation and analysis, and the monitoring of staff competencies. The Safety Case Development Plan becomes a part of the SMS Development Plan. Safety management systems were the core constituent of railway safety cases.

The structure of safety management systems is analogous to that of quality management systems (QMS) (the ISO 9000 series). A good SMS emphasises safety assurance, analogous to quality assurance. Audit plays an important role in both SMS and QMS regimes, and in a well-managed organisation is seen as a driver for continuous improvement. Audit operates at all levels from physical inspection of assets to individual counselling.

There is an equal emphasis on compliance with standards (see below) as a tool for the implementation of continuous improvement.

Standards

It is perhaps not surprising, given the nature of railways, that there is great emphasis on development of, and compliance with, standards. This includes not only hardware and software standards, which adopt the European Committee for Standardization (CEN), European Committee for Electrotechnical Standardization (CENELEC) and International Organization for Standardization (ISO) standards wherever possible, but also operational and competency standards and, at the international level, interoperability standards.

In recent years there has been great emphasis on defining and maintaining staff competencies, with regular assessment and testing, on a par with flight licensing requirements in civil aviation and air traffic control.

All of the above elements have survived the transition from the goal-based safety case regime to the post-2006 certification system implementing the 2004 Rail Safety Directive.

5 Discussion

As with the petrochemical industry analysis, there are two issues of concern: the first is whether or not safety cases have made an improvement to safety performance on the UK's railways. If so, then the next question is the degree to which the type of safety case regime adopted for the UK railway industry from 1994–2006 would be transferable to a healthcare setting. This issue is addressed in the next section.

As described previously, the introduction of the safety case regime onto the UK's railways was driven by and concurrent with privatisation. There has been much debate about whether the fragmentation and consequent communication difficulties due to privatisation were contributory causes of the Southall and Ladbroke Grove disasters and whether rail safety has improved or not since privatisation. One academically rigorous study[9] has suggested that a steady improvement in rail safety, as evidenced by a continuing decline in significant train accidents per unit distance travelled, had continued through the period of privatisation (up to the time of the study, 1994–2003 see Figure F1 below). It is likely that the safety case regime and its emphasis on safety management and risk control had, in fact, contributed to maintaining this trend, despite the negative impact of the fragmentation of British Rail assumed or argued for by many critics, particularly in technical journalism.

Figure F1. UK rail accident statistics 1971–2003[8]

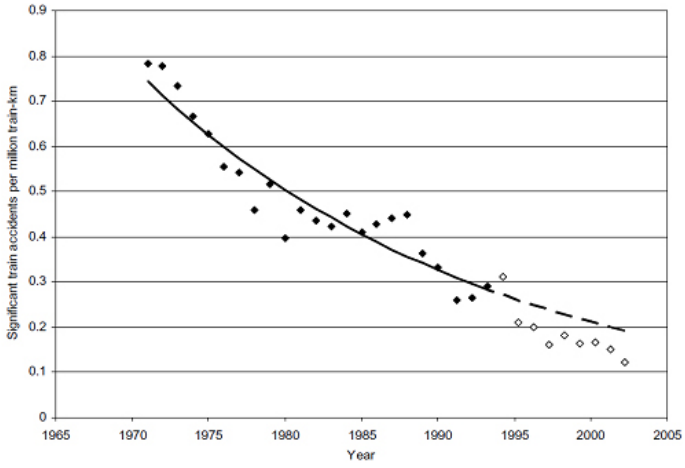

There is little hard evidence to support the proposition that safety cases have made an improvement to safety performance on Britain's railways. We are unaware of any research that demonstrates such a link and consider that it would be virtually impossible to isolate the safety case regime from other potentially relevant factors in the face of the second order change facing the railway industry after 1994. However, the view from inside the industry is that application of the constituent parts of a safety case regime appears to add value. As with other high-risk industries, it would be difficult to argue that reviewing hazards, assessing risks and managing safety could have anything other than a positive impact on safety performance.

It is likely that the UK's railway safety case regime assisted in maintaining safety through the serious upheavals associated with privatisation or part-privatisation. The authors can confirm from personal experience that there are aspects of the quantified risk assessment system on the London Underground (specifically fault and failure monitoring and data collection) that could have been seriously compromised by privatisation if it had not been for the statutory requirement under the R(SC) Regs to collect and consolidate such data in order to maintain the currency of the corporate risk assessment.

6 Lessons and recommendations for healthcare

All of the comments made in Supplement E about the applicability of safety case regimes in high-risk industries to healthcare are applicable, mutatis mutandis, to railways. This section adds a few brief notes that are particular to railways.

Although the safety case regime in the UK railway industry up to 2006 was goal-setting ('permissioning') ra ther than prescriptive, the development of safety cases was dominated by adherence to standards and operating procedures, mostly if not all within the umbrella of the safety management system (SMS). The result was that there was, in reality, perhaps less freedom for individual operators to choose their own analysis methods than in the land-based petrochemical or offshore industries. To that extent, the petrochemical model may be better suited to inform healthcare than the railway model.

However, certain aspects of railway safety management should be of great interest to healthcare. One is the highly structured approach to safety management at all levels, from ensuring that staff wear appropriate protective equipment for hazardous tasks, to evacuating disabled passengers from stations in a serious emergency, such as the terrorist attacks of 7 July 2005.

Another useful model from the railway industry (and aviation) is the emphasis on maintenance of staff competency, particularly of safety critical staff such as train operators, infrastructure inspectors and line station managers. This involves regular testing of both "book" knowledge and the demonstrated ability to carry out key tasks.

The combination of 'hard' QRA techniques and 'soft' human error assessment techniques should also be of interest to the healthcare industry for numerical assessment of risks. The maintenance of statistics in the healthcare industry should place it in a strong position to take advantage of these techniques, as an input to risk control and the avoidance of adverse events.

One of the strengths of the UK rail safety case regime was the system of assuring that the infrastructure controller reviewed and approved the safety cases of its operators. Critics of the system might say that this was merely ensuring an element of cohesion that would not have had to be enforced if the industry had not been fragmented by privatising it in the way that this was done. However, the complexity of modern industrial systems is such that collaboration between different organisations cannot be avoided in the delivery of any substantial service, and systems for communication and sharing of data become essential. The safety case regime ensured that resource was committed to the necessary level of cooperation at all levels.

Above all, the adoption of a safety case regime in the railway industry enabled it to maintain the falling trend in fatality statistics, despite the second order change involved in privatisation. This experience may well benefit, or at least provide encouragement to, the healthcare sector in a time of change.

7 Glossary

See also Supplement E.

Term		Notes
ALARP	As low as reasonably Ppacticable	A concept involving weighing a risk against the trouble, time and money needed to control it.
FMEA	Failure modes and effects analysis	A technique for evaluating the ways in which equipment can fail (or be improperly operated) and the effects these failures can have on a process.
HAZOP	Hazard and operability study	A systematic method for the identification of hazards and problems that might prevent efficient operation of a plant.
RAMS	Reliability, availablity, maintainability and safety studies	A systematic technique for evaluating all aspects of the performance of an asset or system in its operating and maintenance context.
QRA	Quantified risk assessment	Generic term for a range of techniques to assess the numerical probability of adverse events and consequent fatalities, given known or assumed frequencies of occurrence of contributory causes.

8 References

1 Health and Safety Executive, *The tolerability of risk from nuclear power stations*, HSE 1988 revised 1992, Appendix 2, Comparisons of Risk.

2 Health and Safety Commission, *Ensuring Safety on Britain's Railways: A report submitted to the Secretary of State for Transport by the Health and Safety Commission developing proposals for assuring safety following the liberalisation of access to and privatisation of British Railways*. Health and Safety Commission, January 1993.

3 European Parliament and Council Directive 2004/49/EC – *on safety on the Community's railways and amending Council Directive 95/18/EC on the licensing of railway undertakings and Directive 2001/14/EC on the allocation of railway infrastructure capacity and the levying of charges for the use of railway infrastructure and safety certification (Railway Safety Directive)*. Official Journal of the European Union L220/16-39, 21 June 2004.

4 The Railways and Other Guided Transport Systems (Safety) Regulations 2006, UK Statutory Instrument 2006 No.599.

5 EC Directive 91/440/EEC "*on the development of the community's railways*", 29th July 1991]

6 Health and Safety Commission, *Proposals for new safety regulations for railways and other guided transport systems*", Consultative Document, Health and Safety Commission, 2004.

7 BS EN 50126-1:1999 "Railway applications - The specification and demonstration of Reliability, Availability, Maintainability and Safety (RAMS) - Part 1: Basic requirements and generic process

8 Health and Safety Executive, *Railways (Safety Case) Regulations 2000 including 2001 amendments: Guidance on the Regulations*. HSE Books publication L52, 2001 Revision, ISBN 0 7176 2127 8

9 Evans, AW, '*Rail safety and Rail Privatisation in Britain*', Centre for Transport Studies, Imperial College London, 2004.

Supplement G:

Safety case use within the medical devices industry

Robin Bloomfield, Nick Chozos, George Cleland
Adelard LLP

This is one of a series of supplements to the report: *Using safety cases in industry and healthcare: a pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare*

To access the report and the other supplements, please visit www.health.org.uk/safetycasesreport

Contents

Supplement G:

Safety case use within the medical devices industry

1 Introduction

In safety-critical industries such as nuclear energy, aviation and defence, an assurance (or safet y) case is required by the regulator to establish the safety of systems or safety-critical activities. The assurance case is not currently a regulatory requirement in healthcare. However, there have been some recent developments which indicate that the benefits of the assurance case are beginning to be recognised in the medical domain. The most significant development has taken place in the USA, where the regulator of medical devices – the Food and Drug Administration (FDA) – has recently published a guidance document that outlines an assurance case approach for infusion pumps, and it is expected that it will soon become a requirement for approval of these medical devices.

This report contains a review of these recent developments and identifies the challenges that relate to medical devices containing software that have motivated this activity.

Section 2 provides some background regarding medical devices. In particular, it presents an overview of applicable European, UK and USA regulation and focuses on the problems associated with medical devices containing software that have motivated considerations of an assurance case approach.

Section 3 focuses on infusion pumps and the recent developments in the USA, where a requirement for assurance cases has been considered by the regulator, the FDA. In this section, we present an overview of the FDA guidance for assurance cases of infusion pumps.

Section 4 considers other activities that are taking place in the USA and UK that are directly or indirectly related to assurance cases. In particular, we identify ongoing research and publications in this area.

Finally, Section 5 presents a selection of national and international standards that influence the development and safety management of medical devices, and Section 6 provides a report summary.

2 Medical devices

A medical device is any product which is used for medical purposes on patients, in diagnosis, therapy or surgery. There is a wide range of medical devices which play different roles in the delivery of patient care. Increasingly, software is being used to enhance device functionality and improve performance. Despite the benefits, software failure poses a number of additional risks to patient safety; moreover, the complexity of software poses significant challenges for regulators to ensure their safety and confirm that they perform to the manufacturer's specification.

In Europe, medical devices are regulated by the European Commission (EC) Medical Device Directives.[9] These directives place obligations on manufacturers to ensure that their devices are safe and fit for their intended purpose before they are CE marked and placed on the market in any EC member state. In each member state, there is a responsible regulatory authority to implement these EC Directives. In the UK, this authority is the Medicines and Healthcare products Regulatory Agency (MHRA). In the USA, the regulator is the Food and Drug Administration (FDA).

This section provides a background to medical devices and the regulatory framework that attempts to ensure their safety in Europe, the UK and the USA. It then focuses on the problems relating to software-based medical devices and the challenges that regulators face in their evaluation of these devices.

2.1 Definitions of the medical device

A medical device is defined by EC Directive 2007/47/ec[6] as:

> *'any instrument, apparatus, appliance, software, material or other article, whether used alone or in combination, including the software intended by its manufacturer to be used specifically for diagnostic and/or therapeutic purposes and necessary for its proper application, intended by the manufacturer to be used for human beings'.*

In the USA, the FDA provides a similar definition. According to the FDA website, a medical device is:

> *'an instrument, apparatus, implement, machine, contrivance, implant, in vitro reagent, or other similar or related article, including a component part, or accessory which is:*
>
> *– recognized in the official National Formulary, or the United States Pharmacopoeia, or any supplement to them,*
>
> *– intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease, in man or other animals, or*
>
> *– intended to affect the structure or any function of the body of man or other animals, and which does not achieve any of it's primary intended purposes through chemical action within or on the body of man or other animals and which is not dependent upon being metabolized for the achievement of any of its primary intended purposes.'*

There is obviously a very wide range of medical devices. These vary depending on the technology (from syringes to computer software) and the functionality they offer. The EU and USA regulation identifies classes of medical device based on the risks to safety their failure may incur. This classification also facilitates the identification of additional controls that manufacturers have to put in place to ensure regulator approval. The approaches to medical device classification are presented in the following section.

2.2 Classification and regulation

2.2.1 The European approach

The classification of medical devices in the European Union is outlined in Annex IX of the Council Directive 93/42/EEC.[11] There are basically four classes, ranging from low risk to high risk: Class I, Class IIa, Class IIb and Class III.

According to the Directive:

> *'…the classification rules are based on the vulnerability of the human body taking account of the potential risks associated with the technical design and manufacture of the devices; whereas the conformity assessment procedures for Class I devices can be carried out, as a general rule, under the sole responsibility of the manufacturers in view of the low level of vulnerability associated with these products; whereas, for Class IIa devices, the intervention of a notified body should be compulsory at the production stage; whereas, for devices falling within Classes IIb and III which constitute a high risk potential, inspection by a notified body is required with regard to the design and manufacture of the devices; whereas Class III is set aside for the most critical devices for which explicit prior authorization with regard to conformity is required for them to be placed on the market.'*

In order to acquire authorisation of a medical device, the manufacturer has to issue a Declaration of Conformity which must be verified by a Notified Body – a public or private organisation

that has been granted the authority to validate the compliance of the device to the European Directive. In the UK, this role is taken up by the MHRA. In Europe, certified medical devices carry a CE mark.

2.2.2 The USA approach

In the USA, medical devices are regulated by the FDA, and in particular its Center for Devices and Radiological Health (CDRH).[12] The FDA identifies three classes of medical devices, These are based on the level of control necessary to assure the safety and effectiveness of the device.

The following text is an extract from the FDA website, defining the classification system used in the USA:

'Class I: General controls

– Class I devices are subject to the least regulatory control. Class I devices are not intended for use in supporting or sustaining life or to be of substantial importance in preventing impairment to human health, and they may not present a potential unreasonable risk of illness or injury. Examples are elastic bandages, examination gloves, and hand-held surgical instruments.

Class II: General controls with special controls

– In addition to complying with general controls, Class II devices are also subject to special controls. Special controls may include special labelling requirements, mandatory performance standards and post-market surveillance. Devices in Class II are held to a higher level of assurance than Class I devices, and are designed to perform as indicated without causing injury or harm to patient or user. Examples of Class II devices include powered wheelchairs, infusion pumps, and surgical drapes.

Class III: General controls and pre-market approval

– A Class III device is one for which insufficient information exists to assure safety and effectiveness solely through the general or special controls sufficient for Class I or Class II devices. Such a device needs pre-market approval, a scientific review to ensure the device's safety and effectiveness, in addition to the general controls of Class I. Class III devices are usually those that support or sustain human life, are of substantial importance in preventing impairment of human health, or which present a potential, unreasonable risk of illness or injury. Examples of Class III devices which currently require a pre-market notification include implantable pacemaker, pulse generators, HIV diagnostic tests, automated external defibrillators, and endosseous implants.'

The FDA identifies a number of safety regulations and standards that the medical device and the manufacturer must adhere to in order to achieve authorisation. Some requirements only apply to higher-class medical devices.

Table G1: FDA medical device safety regulation

Requirement	Description
Establishment registration	Manufacturers (both domestic and foreign) and initial distributors (importers) of medical devices must register their establishments with the FDA.
Medical device listing	The FDA identifies several types of establishments which must be declared, such as manufacturers, re-packagers, specification developers etc.
Pre-market notification 510(k), unless exempt, or pre-market approval (PMA)	The manufacturer must make a pre-market notification 510(k) submission which demonstrates that the device is 'substantially equivalent to one legally in commercial distribution in the United States'. For Class III devices – high-risk devices – the FDA requires a PMA, which involves a submission of clinical data in support of claims made about the device.
Investigational device exemption (IDE) for clinical studies	An investigational device exemption (IDE) grants permission for the investigational device to be used in a clinical study, which has the purpose to accumulate data on the safety and effectiveness. This data will be used in support a PMA application or a pre-market notification 510(k) submission.
Quality system (QS) regulation	This regulation identifies requirements on the methods, facilities and controls used in the production, installation and servicing of medical devices. The FDA carries out inspections to evaluate whether manufacturing facilities meet these requirements.
Labelling requirements	This covers the labelling on the device and the literature that describes the device and its operation.
Medical device reporting (MDR)	The FDA has a medical device reporting (MDR) programme. When an incident involving a medical device occurs, it must be reported to the FDA and communicated, through this programme, from the manufacturer to the users of the device in a timely manner.

2.3 Medical devices containing software

This report takes a particular focus on medical devices containing software. Software is increasingly used in medical devices as it can substantially increase functionality and effectiveness. Software is found in healthcare in a variety of safety-critical applications, from patient records to radiology devices. However, despite the obvious benefits of software, it also poses new, significant risks to patient safety. Furthermore, the nature of software makes it significantly more difficult for regulators to assess their safety and reliability.

Many concerns have been raised about the design of software in healthcare.[7,10] These concerns focus on problems in the development processes followed by manufacturers, which often result in software errors or inadequate user interfaces. In terms of user interface design, it has been suggested[7] that in many cases, experienced device operators were confused have become lost while configuring devices, have difficulty tracking device states, or they may misinterpret a device function. One example is the case of infusion pumps (Figure G1, overleaf), which are set up to administer drugs using parameters that do not match the way the clinicians think about infusions.

Figure G1: A nurse programming infusion pumps

The complexity of software makes it much harder for regulators to assess these devices and ensure that they provide patients with a high level of protection and attain the functionality and performance levels attributed to them by the manufacturer. Furthermore, it is common practice that software requires the periodic installation of updates or patches; as new code is installed in the device, it is likely that the behaviour of the software also changes, or that unwanted behaviour is introduced, which makes it even more difficult for regulators to confirm the safety of the device once they have been released to the market.

In the USA, the FDA is beginning to consider the use of the assurance case to confront these issues. We have seen that other industries (such as nuclear) take a particular approach when developing assurance cases for devices containing software, with identifiable benefits. The following section considers the recent developments in the USA concerning assurance cases for infusion pumps.

3 Assurance cases and infusion pumps

This section is based on the FDA guidance publication for an assurance case approach for infusion pumps.

3.1 Infusion pumps

External infusion pumps are medical devices that deliver fluids, including nutrients and medications such as antibiotics, chemotherapy drugs, and pain relievers, into a patient's body in controlled amounts.[1] Infusion pumps are often used in combination for a single patient when he or she is being administered a combination of medication (as in Figure G2 below).

Figure G2: Infusion pumps in a neonatal ICU

In general, an infusion pump is operated by a trained user (although incident analyses have identified deficiencies in user training), who programmes the rate and duration of fluid delivery through a software interface built-in on the device. Infusion pumps offer significant advantages over manual administration of fluids. For instance, they offer the ability to deliver fluids in very small volumes, and the ability to deliver fluids at precisely programmed rates or automated intervals.[13]

Because infusion pumps may be used to administer critical fluids, including high-risk medications, pump failures (malfunctions or misuse) can have significant implications for patient safety. As such, many pumps are equipped with safety features. For example, some pumps are designed to alert users when air or some other form of blockage is detected in the tubes that deliver fluid to the patient.

In the past five years (between 1 January 2005 and 31 December 2009), the FDA has seen an alarming increase in the number and severity of infusion pump recalls, while it has received over 56,000 medical device reports (MDRs) associated with the use of infusion pumps. Analyses of these MDRs have revealed device problems that appear to be mainly a result of faulty design. Of these reports, approximately 1 per cent were reported as deaths, 34 per cent were reported as serious injuries, and 62 per cent were reported as malfunctions.

According to the FDA:[13]

> *'…these MDRs identified problems to be attributed to:*
>
> – ***Software Defects.*** *Many of the problems that have been reported are related to software malfunctions. For example, some pumps fail to activate pre-programmed alarms when problems occur, while others activate an alarm in the absence of a problem. Other software errors can lead to over- or under-infusion. In one case, a software problem called a 'key bounce' caused an infusion pump to occasionally register one keystroke (e.g., a single zero, '0') as multiple keystrokes (e.g., a double zero, '00').*
>
> – ***User Interface Issues.*** *There have also been numerous reports of confusing or unclear on-screen user instructions, which may lead to improper programming of medication doses or infusion rates. For example, the design of the infusion pump screen may not make clear which units of measurement (e.g., pounds versus kilograms) should be used to enter patient data, leading to inappropriate dosing.*
>
> – ***Mechanical or Electrical Failures.*** *Other problems that have been reported include components, such as pump housings, that break under routine use; premature battery failures; and sparks or pump fires. Each of these types of incidents can create risks to patients, including the potential for over- or under-administration of critical fluids.'*

In many reports, the manufacturer reported the problem as 'unknown'. Subsequent root cause analyses revealed that the majority of these were design problems and as such foreseeable and, therefore, preventable. Furthermore, these infusion pump problems have been observed across multiple manufacturers, pump types, and use environments.

The increase in the reports of such problems has raised concerns in the FDA. As such, the regulator has begun taking a more proactive and comprehensive approach to prevent safety problems by fostering the development of safer, more effective infusion pumps across the industry. This is done by requiring from the industry the preparation of pre-market notification submissions for infusion pumps and to identify device features that manufactures should address throughout the total product lifecycle. For this, a guidance publication has been developed, which suggests taking an assurance case approach.

3.2 Overview of FDA guidance

The guidance,[2] currently in draft format, is entitled:

– Guidance for Industry and FDA Staff – Total Product Life Cycle: Infusion Pump – Pre-market Notification [510(k)] Submissions

This guidance has been developed in support of pre-market notification submissions for infusion pumps (the requirement for pre-market notification submissions was introduced in Section 2.2.2). The intention of this guidance is to eventually improve the quality of infusion pumps and reduce the number of adverse events and the subsequent device recalls and infusion pump MDRs. FDA's guidance documents are to be viewed only as recommendations, unless specific regulatory or statutory requirements are identified.

However, in the case of infusion pumps, the FDA suggests that:

'…even though the document has been issued as guidance it has been made clear to IP manufacturers that an assurance case is required to obtain pre-market clearance.'

The FDA has committed to a formal pilot programme of use of assurance cases for infusion pumps by the end of March 2011.

3.2.1 Infusion pump assurance case approach

The FDA guidance contains a set of recommendations which are to be documented in the 'pre-market' submission for the approval of an infusion pump. This section provides a summary but focuses more on some of the key elements of the guidance. These are:

– assurance case approach and report (Section 3.2.2)

– risk management activities and report (Section 3.2.3)

The guidance begins by stating the scope of the guidance and providing a definition of the infusion pump system. FDA defines the infusion pump system which includes, apart from the device itself, the 'complete fluid pathway'. That is, the fluid source container, infusion set, extension sets, filters and valves, clamps, up to and including the patient connection and other components. It also includes the description of the patient, the environment of use, and the expected user.

It then outlines the activities and information to be provided.

This guidance is intended to address the root causes of the infusion pump recalls that have been taking place in the past five years. The guidance recommendations are therefore aiming to provide formal, structured methods to a systematic review of the infusion pump system as defined in the beginning of this section.

In order to facilitate this systematic review, the FDA recommends that the activities and information required in Table G2 are also presented in a formal, systematic manner. It is required that a sound approach is taken to demonstrate the validity and completeness of the activities undertaken. For this, the FDA require an assurance case and a report to be submitted as part of the pre-market submission.

Table 9: Overview of the FDA guidance on pre-market notification [510(k)] submissions

Requirement	Description
Device description	The description of the device must clearly indicate the intended use and its environment. The components and features of the device (eg alarms and pump log) must be identified, with particular focus on the user interface. It is also recommended that a comparison of these features with another, already marketed, similar device is provided.
Risks to health	The FDA has identified the risks to health generally associated with the use of infusion pumps. It also recommends a number of measures to mitigate these risks. The guidance suggests that if an alternative approach to address a particular risk is undertaken, there should be sufficient detail to support this decision.
Assurance case report	The FDA describes an assurance case as a formal method for demonstrating the validity of a claim by providing a convincing argument together with supporting evidence. The guidance explains in some detail how the approach to an assurance case should be taken. This is discussed in more detail in Section 3.2.2.
Clinical evaluation	This refers to the evaluation of device performance, possible user error and other human factors issues. Errors identified must be documented. These evaluations are highly recommended for new devices and when there is a major change or modification in the intended use or to correct problems with the design of the user interface.
Risk management	The FDA requires that decisions made in the design and development of the device are documented and reviewed. The results of the risk management activities should be incorporated in the designs. This is also discussed in more detail in Section 3.2.3 as the risk management activities are a key element of the assurance case.
510(k) pre-clearance inspection for infusion pump	The FDA may conduct a pre-clearance inspection for infusion pump manufacturers. Since the recent increase in infusion pump-related MDRs, the FDA carried out a number of inspections which concluded that manufacturers deviated from the quality system regulations. As a result, it was decided that the FDA will conduct pre-clearance inspections to ensure there are no deviations. The FDA provides a guidance manual on the inspection of medical device manufacturers.
Labelling	The FDA specifies labelling requirements for medical devices,[14] which must be adhered to. It is also recommended that the manufacturer provides clear and concise instructions for use that describe the technological features of the infusion pump and how to use the device on patients.
Post-market surveillance of infusion pumps	All manufacturers of medical devices must submit reports to the FDA whenever there is information that a device they market may have caused or contributed to a death or serious injury, or if it has malfunctioned and the malfunction would be likely to cause or contribute to such an incident if it occurred again – this is required by the medical device reporting (MDR) regulation. All of the MDRs can be found in the MDR database.[15]

3.2.2 Assurance case approach and report

One approach to ensuring the safety of medical devices taken by the FDA, is to demonstrate that the device is 'substantially equivalent to one legally in commercial distribution in the United States' (discussed in Table G1). For infusion pumps, the FDA is concerned that due to the 'new implementations of software and changes in materials, design, performance and other features' that most new infusion pumps will contain, it is substantially more challenging for them to assure that equivalence with other marketed devices.

In order to achieve the demonstration of 'substantial equivalence', the FDA recommends that the assurance case approach is undertaken and that an assurance case report is produced as part of the pre-market notification submission.

The FDA draws upon experience in other safety-critical industries (eg nuclear and avionics), believing that 'the methodology will be particularly useful for presenting and reviewing information about infusion pumps'.

3.2.2.1 Assurance case approach

The FDA recommends that manufacturers follow the claims-arguments-evidence (CAE) approach in their assurance case.

These three elements are described in the guidance as:

- **Claim**: statement about a property of the system or some subsystem
- **Evidence**: information that demonstrates the validity of the claim. This can include facts (eg based on observations or established scientific principles), analysis, research conclusions, test data, or expert opinions
- **Argument**: links the evidence to the claim. Arguments can be deterministic, probabilistic, or qualitative. The argument will describe what is being proved or established (ie the claim(s)), identify the items of evidence you are appealing to, and the reasoning (inference, rationale) that the evidence is adequate to satisfy the claim. Arguments may also introduce sub-claims or assumptions which require further exposition.

The FDA guidance continues to focus the assurance case on safety by stating that:

An assurance case addressing safety is called a safety case. A top-level claim (eg 'this infusion pump is comparably safe') is supported by arguments that demonstrate why and how the evidence (eg performance data) supports the top-level claim. The arguments in a safety case are typically organised in a hierarchical fashion with multiple layers of sub-claims, each supported by appropriate evidence. The arguments in a safety case are intended to convince a qualified reviewer or reviewers that the top-level claim is valid.

3.2.2.2 Hazard focus

The FDA has identified a number of known hazards based on the post-market data regarding recalls and MDRs. These hazards are identified in the guidance and it is expected that these will be used as a starting point for the device hazard analysis; these are the minimum set that should be addressed, even though more are likely depending on the device.

Table G3 below lists these hazard categories, although these are discussed in far more detail in the guidance. In particular, the guidance identifies several hazards under each category, their corresponding risks to health and potential causes of the hazard.

Table 10: Hazard categories

Hazard categories	Description
Operational	Hazards inherently related to the operation of the device such as 'air in line' and 'occlusion'.
Environmental	Hazards related to the location in which the device will be used. For instance, 'exposure of the device to hazardous substances' or 'tampering'.
Electrical	Hazards related to the power supply and infrastructure, eg 'overheating' and 'supply voltage error'.
Hardware	Hazards related to the failure of a hardware component of the device, such as 'memory failure' and 'false alarm'.
Software	Hazards related to improper implementation of the development lifecycle for the software. Examples include 'Data error' and 'Software runtime error'. For software, there is also a recommendation to describe how information security will be addressed (confidentiality, integrity, availability and accountability).
Mechanical	Hazards related to the mechanical design of the device by the user, such as 'pump stops infusion' or 'failure to alarm'.
Biological and chemical	Hazards related to the materials of construction, cleaning substances, and infusates. Examples include 'toxicity' and 'loss of drug potency'.
Use	Hazards related to the use of the device and are found within the interaction between the device and the user. For instance, 'the pump is programmed incorrectly' and 'the user fails to understand pump notifications'.

For each of these hazard categories, the guidance identifies control or mitigating activities and additional guidance and standards. Below is an extract from the guidance that focuses on software hazards as an example:

> *'Please refer to the Guidance for the Content of Pre-market Submissions for Software Contained in Medical Devices, http://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm089543.htm, for a discussion of the software documentation that you should provide in the 510(k) submission. We generally consider infusion pumps to be a 'Major' level of concern for the purposes of software review. […]*
>
> *If the device includes off-the-shelf software, you should provide the additional information as recommended in the Guidance for Industry, FDA Reviewers and Compliance on Off-the-Shelf Software Use in Medical Devices,*
>
> *http://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm073778.htm.'*

These hazards should be addressed in the assurance case, although it is expected that the manufacturer will carry out their own hazard analysis.

The pre-market notification should clearly describe the methods used to analyse the hazards and the mitigation of each hazardous event. Alternative approaches to the ones suggested in the guidance are welcome, but there should be sufficient detail provided to support the approach followed.

3.2.3 Risk management

The FDA recommends that the 510(k) submission should also include a risk management report to demonstrate that the results of risk management activities were incorporated into the device design. This report and any supporting documentation should be incorporated into the assurance case as arguments or evidence.

According to the guidance, the following should be included:

'Risk Management Report

– Your submission should include a risk management report, summarizing the results of risk management activities pertaining to safety and effectiveness of the device.

– The risk management report may be submitted in any reasonable format. For example, the information may be presented in tabular or narrative format, and risk levels may be expressed quantitatively or qualitatively as appropriate. The format you choose will probably be dictated by your firm's risk management process.

System Architecture

– Your submission should describe the major components of your system, and indicate how the functional requirements are allocated among them. This may be as simple as a block diagram listing the major functions performed by each system component, and how they are integrated into one system.

Design Requirements Documents

– FDA recommends that submissions include documents resulting from the design input process that define the inputs (i.e., functional, performance, and interface characteristics) of your system in engineering terms. Your firm may refer to these documents as "system specifications", "design requirements", "requirements specifications", or by other names. In many cases, there will be several such documents, covering the major hardware and software components of the system.

– Each submission should also document how the design outputs meet the design inputs. The submission should list the design changes (if any) that were made to the device during the design phase, and how the change affected the device design (i.e., inputs and outputs). The documents provide objective evidence that appropriate risk control measures have been incorporated into the device design, that the device design adequately accounts for the intended use environment, and that technical characteristics of the device critical to clinical performance are adequately described.'

It is anticipated that an FDA investigator will review these same documents during a quality system inspection with the intention to confirm the adequacy of the manufacturer's quality system and the extent to which the procedures are followed.

3.3 Impact of the FDA guidance

Since the publication of this guidance, several 510(k) submissions have been made. However none had been accepted as of October 2010. In fact, several had been rejected after as little as 10 minutes' scrutiny.

As we have seen in the previous section, the FDA has recommended (but does not require) the use of structured safety notations for the development of infusion pump assurance cases. The most common of these are the claims-argument-evidence (CAE) notation or goal-structuring-notation (GSN). Cases may be submitted in hard-copy or softy-copy format, including in bespoke tools supporting these notations.

The FDA has acknowledged that both the industry and the FDA themselves will require further guidance and training on constructing and reviewing assurance cases. The Association for the Advancement of Medical Instrumentation (AAMI) ran a Webinar on assurance and safety cases in September 2010, expecting 20-30 attendees. Around 300 from 70 organisations actually attended.

Adelard ran a course on use of its assurance and safety case environment (ASCE) tool[20] for medical device assurance cases in the USA in September 2010. This course was sold out, with companies sending up to five attendees, including engineers working in areas other than infusion pumps. This was in anticipation of the scope of the assurance case approach widening to other areas.

Recently, the FDA announced that in 2011 there will be a set of 25 actions that will be followed to improve the 'most common path to market' for medical devices. This will also have an impact on the pre-market notification submission requirement for infusion pumps, so it is likely that the guidance on assurance cases may also be subject to revision.

The FDA is considering working with industry and academia to develop one or more 'template' assurance cases for an idealised infusion pump, possibly based on the generic infusion pump model (see section 4.2).

4 Other developments

Apart from the FDA assurance case guidance, there are other developments elsewhere which relate directly or indirectly to assurance cases in healthcare. It should, however, be noted that the developments discussed here are confined to the areas of the development, maintenance and safety assurance of medical devices, but do not consider wider aspects, such as the assurance of a medical service.

4.1 Report from the Software Engineering Institute

The Software Engineering Institute (SEI) of Carnegie Mellon University published a report entitled *Towards an Assurance Case Practice for Medical Devices*[19] in 2009.

The report provides an overview of the safety case approach and goal-structured cases in particular and then focuses on an infusion pump example. It also considers the process of reviewing assurance cases and discusses relevant practices around the assurance case lifecycle, specifically targeted for medical devices.

4.2 Projects

4.2.1 Generic infusion pump model

The University of Pennsylvania (U. Penn), in collaboration with the FDA and the Fraunhofer Center for Experimental Software Engineering (CESE) are developing models of a generic infusion pump.[23] This includes a description of the device, a set of safety requirements, a hazard analysis, and state machine descriptions of its behaviour.

U. Penn is encouraging others to contribute to the development of this work.

4.2.2 The Computer-Human Interaction for Medical Devices (CHI+MED) project

CHI+MED (Computer-Human Interaction for Medical Devices) is an EPSRC-funded project to improve the safety of interactive (programmable) medical devices, such as infusion pumps. The focus is on device design and human factors, and on the identification of ways to reduce the likelihood of medical error.

The project is a six-year inter-disciplinary programme which blends computer science, cognitive psychology and medicine. One of the key aims of the project is to identify interventions that can help manufacturers, clinicians, procurement staff and patients to help reduce the potential human error.

A lot of the work undertaken (eg Vincent, 2010[17]) focuses on the human factors issues concerning infusion pumps and is motivated by the FDA guidance.

4.2.3 Infusion pump safety work at University of Chicago

The Cognitive Technologies Laboratory (CtL) of the University of Chicago is conducting much work on the safety of infusion pumps. The CtL is based within an anaesthesiology department and also takes an interdisciplinary approach with much insight from clinical experience.

There is much focus on user interface design and errors during programming of infusion pumps but also other issues, such as what information should be kept in infusion pump logs and ways of achieving resilience in infusion pump systems.

CtL publications can be found on the laboratory's website.[18]

5 Medical device standards

There are a number of international and national standards that must be followed by medical device manufacturers. Table G4 below lists some of these standards; we have selected three of these as more relevant and discuss them further in this section.

Table G4: Medical device standards

Standard	Description
ANSI/AAMI HE74-2001 – Human factors design process for medical devices	This standard provides guidance on ergonomic and human factors engineering, aiming to improve user and patient safety, system safety and performance, and operator effectiveness in medical device design.
ANSI/AAMI/ISO 13485:2003 – Medical devices	Specifies requirements for a quality management system where an organisation needs to demonstrate its ability to provide medical devices and related services that consistently meet customer requirements and regulatory requirements applicable to medical devices and related services.
ANSI/AAMI/ISO 14971 – Risk management – Part 1: Application of risk management to medical devices	This standard is discussed in more detail in Section 5.1 below.
IEC 60601-1 Ed. 3.0 b – Medical electrical equipment – Part 1: General requirements for basic safety and essential performance	Contains requirements concerning basic safety and essential performance that are generally applicable to medical electrical equipment. For certain types of medical electrical equipment, these requirements are either supplemented or modified by the special requirements of a collateral or particular standard. Where particular standards exist, this standard should not be used alone.
IEC 61508-3 – Functional safety of electrical/electronic/programmable electronic safety-related systems – Part 3: Software requirements	IEC 61508-3:2010 applies to any software forming part of a safety-related system or used to develop a safety-related system.
IEC 62304 Medical device software – Software lifecycle processes	Defines the lifecycle requirements for medical device software. The set of processes, activities, and tasks described in this standard establishes a common framework for medical device software lifecycle processes. Applies to the development and maintenance of medical device software when software is itself a medical device or when software is an embedded or integral part of the final medical device.
IEC 80001 and IEC 80002	These are discussed in more detail in sections 5.2 and 5.3.

5.1 ISO 14971:2007 – Medical device risk management

ISO 14971:2007 defines a risk management process for manufacturers and as such it is highly related to the assurance case approach that the FDA has proposed.

The approach aims to identify the hazards associated with medical devices, to estimate and evaluate the associated risks, to control these risks, and to monitor the effectiveness of the controls. The requirements of ISO 14971:2007 are applicable to all stages of the lifecycle of a medical device.

The requirements of this standard are mandatory in the European Union and the FDA recognises it as an acceptable risk management model. Even in the case where manufacturers choose not to implement ISO 14971, they are still expected to conduct and document a risk management process.[24]

5.2 IEC 80001 – Application of risk management for IT-networks incorporating medical devices

This standard identifies requirements, such as the roles, responsibilities and activities that apply to effective risk management of IT-networks incorporating medical devices. These requirements aim to address safety, effectiveness and data and system security among other issues. The focus of IEC 80001 is on the 'responsible organization' that has purchased a medical device and is incorporating it into an IT network, as well as the manufacturer.

5.3 IEC 80002 – Medical device software

Section 5.1 discussed ISO 14971, an approach to risk management for the development of medical devices. IEC 80002 provides more detailed guidance on the application of ISO 14971 for the development of software in medical devices.

6 Summary

This report presents the recent developments that relate to the application of assurance cases in the area of medical devices. The report focuses on the recent FDA guidance on an approach to assurance cases for pre-market notification submissions of infusion pumps.

The FDA decided to react to the significant rise of MDRs that has been observed in the past five years. They have identified the causes of these to be found in four main areas: software, human factors, hardware design and manufacturing. They are also concerned that the risks associated with infusion pumps are amplified by the high number of infusion pumps available and the frequency of pump use.

In order to address these issues, the FDA has taken up the assurance case approach as part of the pre-market notification submission. The FDA is also proposing that manufacturers take a goal-based approach in their assurance cases, preferably using claims-arguments-evidence (CAE) or goal-structuring-notation (GSN). This has been developed in other industries, such as nuclear and aviation, and was discussed in more detail in Bloomfield et al, 2011.[16]

Since the release of the draft guidance in April 2010, until October 2010, no pre-market submissions had been approved and it is too early to draw any conclusions on the overall experience with the approach.

While there are many generic aspects to assurance cases, we expect that there will be a need for specific medical device issues to be addressed for the approach to succeed. We believe that this will involve: education and training, and tool and methodology development, including the development of further guidance. Finally, for the assurance case approach to succeed, we anticipate the need to resolve various areas of technical difficulty and disagreement highlighted between the different stakeholders.

7 Glossary

Abbreviation	Explanation
AAMI	Association for the Advancement of Medical Instrumentation
ASCE	'Assurance and safety case environment' (Adelard)
CDRH	Center for Devices and Radiological Health
CAE	Claim-arguments-evidence
CtL	Cognitive Technologies Laboratory
FDA	Food and Drug Administration (USA)
IDE	An investigational device exemption
IEC	International Electrotechnical Commission
IP	Infusion pump
MDR	Medical device report/medical device reporting
MHRA	Medicines and Healthcare products Regulatory Agency
PMA	Pre-market approval

8 References

1 Food and Drug Administration, Medical Devices: Infusion Pumps, www.fda.gov/MedicalDevices/ProductsandMedicalProcedures/GeneralHospitalDevicesandSupplies/InfusionPumps/default.htm

2 Food and Drug Administration, Guidance for Industry and FDA Staff – Total Product Life Cycle: Infusion Pump – Premarket Notification [510(K)] submissions, www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm206153.htm#6

3 Food and Drug Administration, FDA news release: FDA to improve most common review path for medical devices, www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm240418.htm

4 Software Quality Research Laboratory, The Pacemaker project, http://sqrl.mcmaster.ca/pacemaker.htm

5 CHI+MED project, www.chi-med.ac.uk/home.php?

6 European Commission, Directive 2007/47/Ec Of The European Parliament And Of The Council, http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2007:247:0021:0055:en:PDF

7 Nemeth C, Nunnally M, O-Connor M, Klock PA, Cook R, Making Information Technology a Team Player in Safety: The Case of Infusion Devices, Advances in Patient Safety, Interface Design for Infusion Devices, Vol. 1, pp. 319-330

8 Leveson N, Turner CS, An investigation of the Therac – 25 accidents, *IEEE Computer*, Vol. 26, No. 7, July 1993, pp. 18-41.

9 European Commission, Medical Device Directives, http://ec.europa.eu/consumers/sectors/medical-devices/documents/guidelines/index_en.htm

10 D Wallace and DR Khun, Failure Modes in Medical Device Software: An Analysis of 15 years of Recall Data, National Institute of Standards and Technology (NIST), http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.154.5890&rep=rep1&type=pdf

11 European Commission, Council Directive 93/42/EEC of 14 June 1993 concerning medical devices, http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31993L0042:en:HTML

12 Food and Drug Administration, Medical Devices, www.fda.gov/MedicalDevices/default.htm

13 Food and Drug Administration, Improvement Initiative, www.fda.gov/medicaldevices/productsandmedicalprocedures/GeneralHospitalDevicesandSupplies/InfusionPumps/ucm205424.htm

14 Food and Drug Administration, Code of Federal Regulations Title 21 – Part 801: Labelling, www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/CFRSearch.cfm?fr=801.109

15 Food and Drug Administration, Medical Device Reports database, www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfmdr/search.cfm

16 Bloomfield R, Chozos N, Cleland G. The Safety Case in the Nuclear Sector, Adelard document reference D/561/13602/1, Version 1.0, February 2011.

17 Vincent C , Mind the gap: What interactive medical device manufacturers need. Interfaces (BCS Interaction specialist group), issue 84, 2010, pp. 14-15.

18 Cognitive Technologies Laboratory, University of Chicago, www.ctlab.org/publications.cfm

19 Software Engineering Institute, Towards an Assurance Case Practice for Medical Devices, Technical note CMU/SEI-2009-TN-018, October 2009.

20 Adelard, Assurance and Safety Case Environment, www.adelard.com/web/hnav/ASCE/choosing-asce/index.html

21 Adelard, Claims, Arguments and Evidence (CAE) assurance case notation, www.adelard.com/web/hnav/ASCE/choosing-asce/cae.html

22 Kelly T, 'Arguing Safety – A Systematic Approach to Safety Case Management', DPhil Thesis YCST99-05, Department of Computer Science, University of York, UK, 1998.

23 University of Pennsylvania. The Generic Infusion Pump (GIP), http://rtg.cis.upenn.edu/gip.php3

24 eHow.com, What is ISO 14971, www.ehow.com/about_6578424_iso-14971_.html

25 Software Engineering Process Technology, Medical Device Standards, www.12207.com/Medical%20standard.htm

Supplement H:

Safety case use in healthcare – screening results of the literature survey

Mark-Alexander Sujan
Warwick Medical School

This is one of a series of supplements to the report: *Using safety cases in industry and healthcare: a pragmatic review of the use of safety cases in safety-critical industries – lessons and prerequisites for their application in healthcare*

To access the report and the other supplements, please visit www.health.org.uk/safetycasesreport

Contents

Supplement H:

Safety case use in healthcare – screening results of the literature survey

1 Introduction

This chapter describes the design and the results of the systematic review of published literature for evidence of the purposeful adoption of safety cases in healthcare carried out between January and March 2011. The systematic review was complemented by hand searches and expert input.

2 Scope

The aim of the systematic literature review was to identify publications that describe the purposeful application of safety cases (assurance cases) within healthcare. Included were theoretical contributions and opinions as well as empirical studies. Excluded was literature describing the application of safety cases in other industries, theoretical developments of safety case development and maintenance approaches without reference to an application domain, publications that deal primarily with hazard and risk analysis methodologies rather than the demonstration that a system is acceptably safe, and publications relating to relevant standards such as ISO 14971, IEC 60601, ISO 80001, etc, when the focus was on the content of the standard rather than on the demonstration of safety. Excluded were also studies that deal with health and safety aspects of, for example, nuclear waste materials.

3 Search strategy

The search strategy was designed to identify publications in English in journals and published conference proceedings. Key health and non-health databases were searched. To complement the systematic search of published literature, a selective web-based search was performed to identify additional reports, standards and other relevant documents. Finally, expert opinions and views were sought through personal communication.

The following databases were included in the systematic search.

- Health databases: Medline, Health Business Elite, NHS Evidence, NHS Evidence – innovation and improvement, NHS evidence – health management.
- Computer science/engineering databases: ABI Inform Global, Compendex, Inspec, CSA Technology Research Database, Springer Lecture Notes in Computer Science.

The search employed the following strategy:

Concepts and limits	Terms
Safety case	#1 "safety case" OR "safety cases" #2 "assurance case" OR "assurance cases" #3 "safety argument" OR "safety arguments" #4 "trust case" OR "trust cases" #5 "risk management file" OR "risk management files" #6 "integration profile" OR "integration profiles" #7 "interoperability testing" #8 (#1 OR #2 OR #3 OR #4 OR #5 OR #6 OR #7)
Health	*n.b. only necessary in non-health databases* "health care" OR medical OR medicine OR patient* OR hospital* OR ambulance* OR "emergency services" OR "quality of care" OR "health informatics" OR biomedical
Date	No further back than 2000.

4 Searches performed

The following searches were performed:

Medline

Medline (1950 onwards) was searched using the OvidSP interface on 09/02/2011 for the period 01/01/2000 to 10/02/2011.

1. (safety case or safety cases).mp. (34)
2. (assurance case or assurance cases).mp. (4)
3. (safety argument or safety arguments).mp. (4)
4. (trust case or trust cases).mp. (127)
5. (risk management file or risk management files).mp. (8)
6. (integration profile or integration profiles).mp. (55)
7. (interoperability test or interoperability tests or interoperability testing).mp. (3)
8. 1 or 2 or 3 or 4 or 5 or 6 or 7 (232)
9. 8 not (Wellcome Trust Case-Control Consortium or Wellcome Trust Case Control Consortium).mp. (122)
10. 9 (122)
11. limit 10 to yr="2000 -Current" (108)

[mp=protocol supplementary concept, rare disease supplementary concept, title, original title, abstract, name of substance word, subject heading word, unique identifier]

ABI/INFORM Global

ABI/INFORM Global was searched using the Proquest interface on 10/02/2011 for the period 01/01/2000 to 10/02/2011.

1. (("safety case*" OR "assurance case*" OR "safety argument*" OR "trust case*" OR "risk management file*" OR "integration profile*" OR "interoperability test*") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR "quality of care" OR "health informatics")) in Citation and abstracts (39)

[Fewer 'healthcare' terms were used due to the limited number of characters allowed in the search box in this interface.]

Compendex

Compendex was searched using the Engineering Village interface on 10/02/2011 for the period 01/01/2000 to 10/02/2011.

1. ("safety case" OR "safety cases" OR "assurance case" or "assurance cases" OR "safety argument" OR "safety arguments" OR "trust case" OR "trust cases" OR "risk management file" OR "risk management files" OR "integration profile" OR "integration profiles" OR "interoperability test" OR "interoperability tests" OR "interoperability testing") wn KY AND ("health care" OR medical OR medicine OR patient* OR hospital* OR ambulance* OR "emergency services" OR "quality of care" OR "health informatics" OR biomedical) wn KY

(48)

[wn KY = Retrieves results from fields, including abstract, title, translated title, controlled terms, and uncontrolled terms.]

Inspec

Inspec was searched using the Engineering Village interface on 10/02/2011 for the period 01/01/2000 to 10/02/2011.

1. ("safety case" OR "safety cases" OR "assurance case" or "assurance cases" OR "safety argument" OR "safety arguments" OR "trust case" OR "trust cases" OR "risk management file" OR "risk management files" OR "integration profile" OR "integration profiles" OR "interoperability test" OR "interoperability tests" OR "interoperability testing") wn KY AND ("health care" OR medical OR medicine OR patient* OR hospital* OR ambulance* OR "emergency services" OR "quality of care" OR "health informatics" OR biomedical) wn KY

(55)

[wn KY = Retrieves results from fields, including abstract, title, translated title, controlled terms, and uncontrolled terms.]

Health Business Elite

Health Business Elite was searched using the Ebsco interface on 17/02/2011 for the period 2000 to 2011.

1. ("safety case" OR "safety cases" OR "assurance case" OR "assurance cases" OR "safety argument" OR "safety arguments" OR "trust case" OR "trust cases" OR "risk management file" OR "risk management files" OR "integration profile" OR "integration profiles" OR "interoperability test" OR "interoperability tests" OR "interoperability testing").ti,ab,sh [Limit to: Publication Year 2000-2011]

(47)

[ti,ab,sh = title, subject headings, abstract]

NHS Evidence (including 'NHS Evidence – innovation and improvement' and 'NHS Evidence – health management')

NHS Evidence, which includes 'NHS Evidence – innovation and improvement' and 'NHS Evidence – health management', was searched on 17/02/2011 with no date limits.

1. "safety case" OR "safety cases" OR "assurance case" OR "assurance cases" OR "safety argument" OR "safety arguments" OR "trust case" OR "trust cases" OR "risk management file" OR "risk management files" OR "integration profile" OR "integration profiles" OR "interoperability test" OR "interoperability tests" OR "interoperability testing"

(212)

Lecture Notes in Computer Science

Lecture Notes in Computer Science was searched using the SpringerLink interface on 24/02/2011 with no date limits.

1. ("safety case" OR "safety cases") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR "quality of care" OR "health informatics")) AND pub:("Lecture Notes in Computer Science")

(29)

2. ("assurance case" OR "assurance cases") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR "quality of care" OR "health informatics")) AND pub:("Lecture Notes in Computer Science")
(8)

3. ("safety argument" OR "safety arguments") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR "quality of care" OR "health informatics")) AND pub:("Lecture Notes in Computer Science") (17)

4. ("trust case" OR "trust cases") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR "quality of care" OR "health informatics")) AND pub:("Lecture Notes in Computer Science")
(5)

5. ("risk management file" OR "risk management files") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR "quality of care" OR "health informatics")) AND pub:("Lecture Notes in Computer Science")
(1)

6. ("integration profile" OR "integration profiles") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR "quality of care" OR "health informatics")) AND pub:("Lecture Notes in Computer Science")
(8)

7. ("interoperability test" OR "interoperability tests" OR "interoperability testing") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR "quality of care" OR "health informatics"))
(43)

[Due to limitations of this interface (eg the limited number of characters allowed in the search box, no feature to enable the combining of sets using Boolean), 'safety case' terms had to be combined with 'healthcare' terms in several small sets.]

[Fewer 'healthcare' terms were used due to the limited number of characters allowed in the search box in this interface.]

CSA technology research database

CSA technology research database was searched using the CSA Illumina interface on 24/02/2011 for the period 2001 to 24/02/2011.

1. ("safety case" OR "safety cases" OR "assurance case" or "assurance cases" OR "safety argument" OR "safety arguments" OR "trust case" OR "trust cases" OR "risk management file" OR "risk management files" OR "integration profile" OR "integration profiles" OR "interoperability test" OR "interoperability tests" OR "interoperability testing") AND ("health care" OR medical OR medicine OR patient* OR hospital* OR ambulance* OR "emergency services" OR "quality of care" OR "health informatics" OR biomedical)
(47)

Screening results

In total, 667 publications were returned (duplicates not removed). Medline, ABI Inform Global, Compendex and Inspec were screened using titles first (n=210), then on relevant abstracts (n=97) to identify 10 papers for inclusion in the review. Health Business Elite, NHS Evidence (including innovation and improvement, and health management) and CSA Technology Research Database were screened on abstracts (n=300) to identify two papers. Springer Lecture Notes in Computer Science were screened on abstracts (n=111) and eight papers were identified for inclusion. Hand searches identified a further two publications. After removal of duplicates, 16 papers were selected for review (see Table H1).

Three broad application domains were identified and the papers grouped accordingly.

– **Medical devices:** As medical devices increasingly contain programmable elements that make them more flexible, adaptable and allow greater interconnectivity, the assurance that safety objectives are met and the certification of medical devices are also becoming more difficult. This has given rise to new developments reflected in standards and guidance issued by bodies such as the US Food and Drug Administration (FDA). The design and implementation of such programmable medical devices may benefit from methods and techniques developed within the

software and engineering communities for the development of safety-critical computer-based systems, including the use of safety cases.

- **Health informatics:** Within healthcare, there are also an increasing number of health informatics products that are not strictly speaking medical devices, but which may still present risks to the patient. Examples are Computer physician order entry (CPOE) systems or clinical decision support systems.
- **Health systems:** Finally, as the discipline of patient safety matures and organisations are gaining experiences and expertise in the application of systematic methods for identifying and managing risks to patient safety, the use of safety cases or structured safety arguments may be a useful way of documenting and guiding an organisation's safety efforts.

In addition, a selection of relevant standards, ISO 14971, IEC 60601-1, IEC 80001-1, DSCN 14/2009, DSCN 18/2009 and the Care Quality Commission (CQC) Essential Standards of Quality and Safety were reviewed.

6 Results

Medical devices

The majority of papers (9/16) dealt with the application of safety cases for medical devices. This is an expected result since some medical devices pose serious recognised risks to patients (eg radiation therapy machines, infusion pumps etc), and technologically sophisticated devices are close to the engineering communities that have applied safety cases in other safety-critical domains.

Standards review

There is a large number of safety-related standards for medical devices. Two key standards are ISO 14971 (The Application of Risk Management to Medical Devices) and IEC 60601-1-1 (Medical Electrical Equipment – Part 1: General Requirements for Basic Safety and Essential Performance). ISO 14971 requires that the manufacturer shall establish, document and maintain throughout the lifecycle 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. To this end, the manufacturer needs to produce and maintain a risk management file. IEC 60601-1-1 lays out general requirements for safety, while a number of collateral standards IEC 60601-1-x address particular aspects, such as electromagnetic compatibility, usability, etc. IEC 60601-2-x is a series of standards that define safety requirements for specific systems. IEC 60601-1-1 incorporates a risk management lifecycle as defined by ISO 14971.

Literature review

All of the papers engage with the above standards to a certain extent and emphasise different aspects such as implications for device manufacturers, demonstration of compliance with standards, and limitations of focusing on medical devices in isolation.

A number of papers explain the relevant standards to the device manufacturer community and make recommendations for compliance.[1,3,10] There are many more papers of this sort, but the papers that were included in this review also put an explicit emphasis on the demonstration of safety through a structured argument (even if it is not referred to as a safety case explicitly).

Cyra and Gorski[4] provide guidance and develop a template for demonstrating compliance with ISO 14971. This is structured in the form of claims-argument-evidence, which is typical for safety arguments.

A set of three papers from a recent conference on the application of formal methods argue for the benefit of building safety cases from formal models or for the integration of evidence derived from the application of formal methods.[7,8,9]

Finally, two papers address the importance of considering the actual context of use of medical devices. Sujan *et al*[16] argue for an operational safety case to be produced by the healthcare organisation operating the medical device with support by the manufacturer who in turn should produce a generic device safety case. Becker[2] recognises the importance of the context of use, but recommends that the manufacturer should remain in control of relevant information and should instead produce a number of use case safety cases.

Health informatics

Six papers identified during the review dealt with the broader notion of health informatics applications. Some of the papers and standards deal, strictly speaking, with networks of medical devices, but were included in this category rather than in the medical device category since they take a systems view going beyond the individual medical device.

Standards review

IEC 80001-1 (Application of risk management for IT networks incorporating medical devices – Part1: Roles, responsibilities and activities) was published in 2010 and proposes a risk management lifecycle based on ISO 14971 for networks of medical devices. The standard addresses those concerns that were raised previously about the lack of consideration of the integration and interoperability of different medical devices into a system within a healthcare organisation. It defines roles and responsibilities both on the part of the manufacturer as well as on the part of the user organisation. Like ISO 14971, it requires documentation of risk management activities in the risk management file.

The two standards DSCN 14/2009 (Application of clinical risk management to the manufacture of health software) and DSCN 18/2009 (Guidance on the management of clinical risk relating to the deployment and use of health software) were developed as part of the National Programme for IT in parallel with the work on IEC 80001, but they are wider in their scope (considering health informatics applications in general rather than networks of medical devices) and more explicit in their recommendation in the adoption of the safety case concept on the part of both the manufacturer of health software and the healthcare organisation that uses these. While the standards are oriented as before on ISO 14971, they define and specify the notion of a clinical safety case and make explicit the interface between manufacturers and users. Connecting for Health reports positive experiences with the safety case approach, but points out that education and support are vital and that there are still problems with the uptake by manufacturers.[17]

Literature review

The papers identified deal with decision support systems, networks of medical devices and a large-scale drug management IT solution. The papers emphasise the systems aspects of technology, reinforcing the fact that safety assurance needs to account properly for the context of use and the possible interactions that may result. Safety cases are proposed to facilitate communication and to document assumptions.

Two papers[12,13] comment on the development of IEC 80001 as the transition from risk management for individual devices towards risk management for networks of medical devices is made.

The remaining four papers all look at health informatics applications other than medical devices in the strict sense. An early paper[5] argues that the safety aspects of health informatics applications have not been properly addressed by the community, and the paper proposes a structured approach to safety assurance based on methods from artificial intelligence. Another early paper[14] reviews part of the safety evidence for a decision support system in mammography and finds that stronger safety arguments could be created by representing proposed safety barriers (controls or risk reduction measures) explicitly in the safety arguments. Two larger case studies are reported in Gorski *et al*, 2005[6] and Salvaneshi, 2010[11] – albeit with little actual detail. Gorski et al[6] describe the Trust Case development for a health informatics product developed as part of an EU-funded project (Drug in Virtual Enterprise – Drive). The paper finds that there are many assumptions in the safety evidence which should be properly identified and documented as such. Salvaneschi[11] describes the safety testing of a computer physician order entry (CPOE) system in an Italian hospital and proposes Petri nets as a method for deriving test cases that should be used in providing safety evidence.

Health systems

The final category considers health systems in general, including those aspects that are not concerned with the assurance of patient safety related to IT systems. Only one position paper was identified in this category. This is not surprising, since the safety case concept has been developed within the engineering communities and is traditionally applied to hardware and software products first. As the safety case concept takes hold in the domains of medical devices and health informatics applications, we may expect to see more work also in this category of general health systems in the future.

Standards review

The Care Quality Commission (CQC) Essential Standards of Quality and Safety consist of 28 regulations (and associated outcomes) that are set out in two pieces of legislation: the Health and Social Care Act 2008 (Regulated Activities) Regulations 2010 and the Care Quality Commission (Registration) Regulations 2009 (CQC, Quick Guide). Sixteen of these regulations relate directly to safety and quality of care. Each regulation defines an expected outcome or goal that patients can expect in relation to the care that they receive. The standards thus are goal-based, similar to the structures adopted in safety cases. Compliance with the regulations is assessed in a number of ways. Important indicators are the Quality and Risk Profiles (QPR) that are derived from a number of triggers for each regulation and each healthcare organisation.

Literature review

The paper identified during the literature search[15] is a position paper that describes the disjoint regulation between manufacturers of medical devices and healthcare organisations as providers of healthcare services. The paper outlines, using two small examples, how safety cases could be useful to identify and to document emergent interactions resulting from changes in healthcare systems. The examples provided relate to the introduction of bar-coding for patient identification (a technological example) and the introduction of a urinary tract infection (UTI) specialist nurse (non-technological example).

7 Further activities

EWICS Assurance Case Pre-Standardisation Work

The Medical Devices Subgroup of the European Workshop on Industrial Computer Systems (EWICS, www.ewics.org) is promoting the development of safety/trust/assurance cases in the healthcare domain by means of its international workshops, standards review, papers and the development of safety case examples (medical bed, networked systems). The recent standard IEC 80001-1 Ed.1 (2010) on application of risk management for IT-networks incorporating medical devices is an important, but small step forward, acknowledging the need for risk management in network configurations of medical devices. The EWICS Medical Devices and Security Subgroups together are working on a generic case study on the safety and security of wireless medical sensor networks. The discussion and further exploration of this kind of assurance case development will also be the theme for an international workshop that EWICS organises in January 2011 in either Germany or the UK. Special focus is on the relevance of use-in-context expertise of users and how this can be valued in earlier lifecycle stages of new generations of medical devices.

University of Florida

Prof Robert Wears (Department of Emergency Medicine, University of Florida) was leading work on the development of a safety case for the computerisation of an Emergency Department status board. The project did not produce a detailed safety case due to the nature of the project and the fact that actual safety evidence was scarce.

8 Discussion

The systematic literature review demonstrated that research on and application of safety cases to healthcare is scarce. The majority of papers identified described different aspects relating to safety assurance of medical devices. Within the standardisation community there is currently a lively debate around these issues, and it appears that developments are driven by these efforts. This extends to aspects of networked medical devices, where a key standard is the main focus and driver for developments in this direction. It is not clear to what extent manufacturers of medical devices are actively supporting the adoption of the safety case concept.

There were also some examples where the safety case concept has been applied to the wider health informatics field. Despite encouraging findings from, for example, Connecting for Health, there appears to be little awareness of these developments within the health informatics or patient safety community.

Apart from a position paper, there is no evidence that safety cases have been applied to the wider health system where the focus has not been on the introduction of technology.

The literature review suggests that the main drivers for developments currently are the standardisation efforts of organisations such as the US Food and Drug Administration. This appears to be an important factor in securing the attention of the industry. The literature review further suggests that healthcare organisations need to take greater responsibility for actively compiling evidence that the complex systems they operate to provide patient care are, in fact, safe. This, however, will only be possible when adequate resources and training opportunities are provided to these organisations to enable them to build up the capability that is required.

Table H1: Classification of papers identified from the systematic literature search

Domain	Authors	Type	Comments
Medical devices	Bartoo (2003)	Position paper	Discussion of risk management according to ISO 14971. Makes recommendations about systematic documentation of risk management activities in the risk management file.
	Becker (2008)	Empirical / case study	Emphasises the importance of considering the use context in safety assessments of medical devices. Presents a real case study (development of an intensive care work station system) with little actual detail.
	Bills (2006)	Position paper	Describes risk management within IEC 60601-1 and outlines implications for manufacturers of biomedical instruments.
	Cyra, Gorski (2007)	Theory	Development of a template for demonstrating compliance with ISO 14971.
	Huhn, Zechner (2010)	Theory	Development of a development / implementation activity-based quality model relating facts about the software product to development activities to support compliance with IEC 62304.

Domain	Authors	Type	Comments
	Jee, Lee, Sokolsky (2010)	Empirical / case study	Description of model-based implementation of time-critical software for a pacemaker. The assurance case is developed in outline based on the argument that the requirements are consistent and have been met in the implementation (assuming that a hazard and risk analysis has been done previously).
	Lawford, Maibaum, Wassyng (2010)	Position paper	Argues that the current focus with certification of medical devices is on the development process and that product characteristics need to be considered to a greater extent. Safety cases / assurance cases could be useful to introduce objective criteria rather than expert opinion when the development process follows a model-based approach.
	Lincoln (2009)	Position paper	Explains risk management under ISO 14971 and emphasises that the documentation produced in the risk management report can also serve useful internal purposes other than regulation.
	Sujan, Koornneef, Voges (2007)	Position paper	Extends the concept of safety cases for medical devices to include also operational safety cases that demonstrate the safety of medical devices in actual operation. Emphasises the active role of the user of medical devices using the example of medical beds.
Health informatics	Gorski *et al* (2005)	Empirical / case study	Describes the reasoning behind the development of a trust case for a health informatics product developed as part of an EU-funded project (Drug in Virtual Enterprise – Drive). The focus of the paper is on the description of the modelling language, little detail is provided on the safety case. Reports that 97 claims were supported by 234 facts and 121 assumptions. Argues that a trust case is useful to document where assumptions are used to support claims.
	Fox (2001)	Theory	Early paper arguing that safety assurance of health informatics products such as clinical decision support systems have not been safety assured properly. Makes the case for an approach to safety assurance based on artificial intelligence.

Domain	Authors	Type	Comments
	Salvaneschi (2010)	Empirical / case study	Case study describing the safety testing of a computer physician order entry (CPOE) system in an Italian hospital. Argues that clinical information systems have not been properly safety assured. Proposes a testing approach where test cases are developed from a petri-net representation. This forms part of the safety case that should be produced for such products.
	Schrenker (2008)	Position paper	Reviews ISO 14971 and argues the need for applying safety cases to networked systems. Provides a rationale and explains what this entails in practice with a view to IEC 80001.
	Schrenker (2010)	Position paper	Reviews ISO 14971 and argues the need for considering networked systems. Refers to IEC 80001 and makes a case for the role of clinical engineer to ensure consideration of the application context beyond individual medical devices.
	Smith, Harrison, Schupp (2004)	Theory	Identifies different types of barriers in hazard mitigation argument in two small case studies and argues that stronger safety arguments can be constructed by representing these barriers explicitly.
Health system	Sujan *et al* (2006)	Position paper	Argues that safety cases could be a useful tool to identify possible interactions when introducing changes to healthcare systems (both technical and non-technical) – eg introduction of bar-coding for patient identification or introduction of a new role such as UTI specialist nurse.

9 References

1 Bartoo G. Risk management. *IEEE Engineering in Medicine & Biology Magazine*, 2003;22(4):166-70.

2 Becker U. Applying safety goals to a new intensive care workstation system. In Harrison MD & Sujan MA (eds.) *Computer Safety, Reliability & Security* (Safecomp 2008), LNCS 5219;263-276, Springer Verlag.

3 Bills E. Cleaning, sterilization, and biocompatibility: Risk management for IEC 60601-1 third edition. *Biomedical Instrumentation and Technology*, 2006;40;390-2.

4 Cyra L & Gorski J. Supporting compliance with safety standards by trust case templates. In Aven T & Finnem JT (eds.), *Risk, Reliability & Societal Safety (Vol 2)*, pp. 1367-1374, Taylor & Francis, 2007.

5 Fox J. Designing safety into medical decisions and processes. In Voges U (ed.), *Computer Safety, Reliability & Security* (Safecomp 2001), LNCS 2187;1-13, Springer Verlag.

6 Gorski J, Jarzebowicz A, Leszczyna R, Miler J, Olszewski M. Trust case: justifying trust in an IT solution. *Reliability Engineering and System Safety*, 2005;89(1):33-47.

7 Huhn M & A Zechner. Arguing for software quality in an IEC 62304 compliant development process. In *ISoLA 2010*, LNCS 6416;296-311, Springer Verlag.

8 Jee E, Lee I, Sokolsky O. Assurance cases in model-driven development of the pacemaker software. In *ISoLA 2010*, LNCS 6416;243-356, Springer Verlag.

9 Lawford M, Maibaum T, Wassyng A, editors. Certification of software-driven medical devices. In *ISoLA 2010*, LNCS 6416;293-295, Springer Verlag.

10 Lincoln J. Product Risk Management Under ISO 14971:2007. *Journal of Validation Technology*, 2009;15(4):10.

11 Salvaneschi P, editor. Safety testing of computerized provider order entry systems. In *SEHC 2010*;74-82, ACM.

12 Schrenker R. Sufficient Evidence: Making the Case for Safety. *Biomedical Instrumentation & Technology* 2008; 42(6):471-3.

13 Schrenker R. The case for a systems focus in healthcare. *Biomedical Instrumentation & Technology* 2010; Suppl 13-7.

14 Smith SP, Harrison MD, Schupp BA. How explicit are the barriers to failure in safety arguments? In Heisel M *et al* (eds), *Computer Safety, Reliability, and Security* (Safecomp 2004), LNCS 3219;325-337, Springer-Verlag.

15 Sujan MA, Harrison MD, Steven A, Pearson PH, Vernon SJ. Demonstration of safety in healthcare organisations. In Gorski J (ed), *Computer Safety, Reliability, and Security* (Safecomp 2006), LNCS 4166;219-232, Springer Verlag.

16 Sujan M-A, Koornneef F, Voges U. Goal-based safety cases for medical devices: Opportunities and challenges. In Saglietti F & Oster N (ed), *Computer Safety, Reliability, and Security* (Safecomp 2007), LNCS 4680;14-27, Springer Verlag.

17 Baker M, Harrison S. Lessons from NHS Connecting for Health, 2011. Presentation available from www2.warwick.ac.uk/fac/med/staff/sujan/research/safety_case_review/wp3_workshop/baker_harrison_scr.pdf