IChemE ADVANCING CHEMICAL ENGINEERING WORLDWIDE

Fifth Edition

THE SAFETY CRITICAL SYSTEMS HANDBOOK

A STRAIGHTFORWARD GUIDE TO FUNCTIONAL SAFETY:
IEC 61508 (2010 EDITION), IEC 61511 (2016 EDITION)
ALSO RELATED GUIDANCE ON CYBER SECURITY
& INCLUDING MACHINERY AND OTHER INDUSTRIAL SECTORS

Dr. David J. Smith
and Kenneth G.L. Simpson

B/H

The Safety Critical Systems Handbook

A Straightforward Guide To Functional Safety:
IEC 61508 (2010 Edition), IEC 61511 (2016 Edition)
Also Related Guidance on Cyber Security
& Including Machinery and Other Industrial Sectors

FIFTH EDITION

Dr. David J. Smith

Kenneth G.L. Simpson

Butterworth-Heinemann
An imprint of Elsevier

Butterworth-Heinemann is an imprint of Elsevier
The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, United Kingdom
50 Hampshire Street, 5th Floor, Cambridge, MA 02139, United States

Notices

Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

Library of Congress Cataloging-in-Publication Data
A catalog record for this book is available from the Library of Congress

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library

ISBN: 978-0-12-820700-0

Publisher: Susan Dennis
Acquisition Editor: Anita Koch
Editorial Project Manager: Charlotte Kent
Production Project Manager: Prem Kumar Kaliamoorthi
Cover Designer: Alan Studholme

Typeset by TNQ Technologies

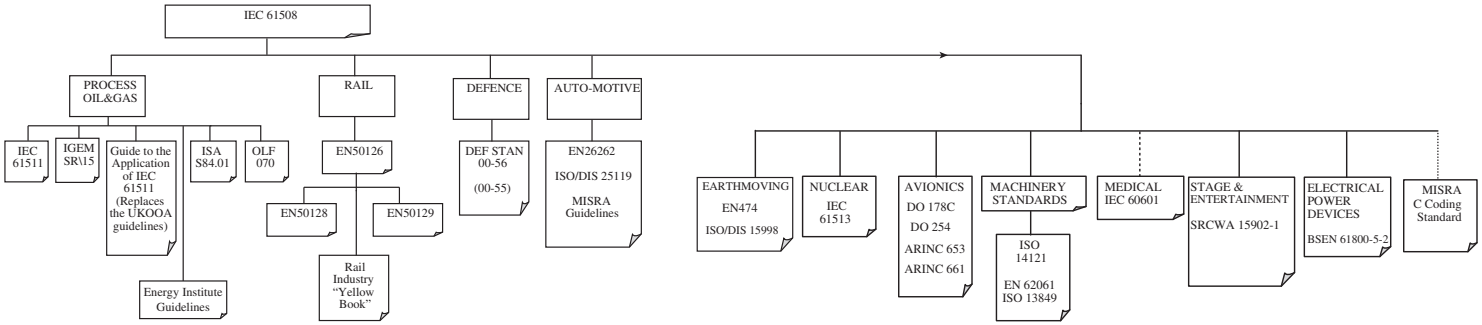
IEC 61508
PROCESS OIL&GAS
RAIL
DEFENCE
AUTO-MOTIVE
IEC 61511
IGEM SR\15
Guide to the Application of IEC 61511 (Replaces the UKOOA guidelines)
ISA S84.01
OLF 070
Energy Institute Guidelines
EN50126
EN50128
EN50129
Rail Industry "Yellow Book"
DEF STAN 00-56
(00-55)
EN26262
ISO/DIS 25119
MISRA Guidelines
EARTHMOVING
EN474
ISO/DIS 15998
NUCLEAR
IEC 61513
AVIONICS
DO 178C
DO 254
ARINC 653
ARINC 661
MACHINERY STANDARDS
ISO 14121
EN 62061
ISO 13849
MEDICAL
IEC 60601
STAGE & ENTERTAINMENT
SRCWA 15902-1
ELECTRICAL POWER DEVICES
BSEN 61800-5-2
MISRA C Coding Standard

A Quick Overview

Functional safety engineering involves identifying specific hazardous failures which lead to serious consequences (e.g., single or multiple deaths, environmental damage) and then establishing maximum tolerable frequency targets for each mode of failure. Equipment whose failure contributes to each of these hazards is identified and usually referred to as "safety-related". Examples are industrial process control systems, process shut down systems, rail signalling equipment, automotive controls, medical treatment equipment, etc. In other words, any equipment (with or without software) whose failure can contribute to a hazard is likely to be safety-related.

A safety-function is thus defined as a function, of a piece of equipment, that maintains it in a safe state, or brings it to a safe state, in respect of some particular hazard.

Since the publication of the first four editions of this book, in 2001, 2004, 2011, and 2016, the application of IEC 61508 has spread rapidly through most sectors of industry. Also, the process sector IEC 61511 document was updated and reissued in 2016. IEC 61508 (BS EN 61508 in the UK) was reissued in 2010. The opportunity has now been taken to update and enhance this book in the light of the authors' further experience. There are still three chapters on industry sectors (Chapters 8, 9 and 10) and Chapters 15 and 16 provide even more examples. Chapter 17 has been added to address the topic of cyber security which is of growing importance and is now called for by many documents.

There are both **random hardware failures** which can be quantified and assessed in terms of failure rates, and **systematic failures** which cannot be quantified. Therefore, it is necessary to have the concept of integrity levels so that the systematic failures can be addressed by levels of rigor in respect of the design techniques and operating activities.

The maximum tolerable failure rate that we set, for each hazard, will lead us to an integrity target for each piece of equipment, depending upon its relative contribution to the hazard in question. These integrity targets, as well as providing a numerical target to be met, are also expressed as "safety-integrity levels" according to the severity of the numerical target. This usually involves four discrete bands of "rigor" and is explained in Chapters 1 & 2.

SIL 4: the highest target and most onerous to achieve, requiring state of the art techniques (usually avoided)
SIL 3: less onerous than SIL 4 but still requiring the use of sophisticated design techniques
SIL 2: requiring good design and operating practice to a level such as would be found in an ISO 9001 management system
SIL 1: the minimum level but still implying good design practice
<SIL 1: referred to (in IEC 61508 & other documents) as "not−safety related" in terms of compliance. (A misnomer in that a quantitative target will exist, which needs to be met, albeit at less than the boundary for SIL 1)

An assessment of the design, of the designer's organisation and management, of the operator's and the maintainer's competence and training should then be carried out in order to determine if the proposed (or existing) equipment actually meets the target SIL in question.

Overall, the steps involve

Setting the SIL targets—Chapter 2.1
Capability to design for functional safety—Chapter 2.2
Quantitative assessment—Chapters 3, 4, 5 & 6
Qualitative assessment—Chapters 3 4 & 17
Establishing competency—Chapter 2.3
As low as reasonably practicable—Chapter 2.2 & 2.4
Reviewing the assessment itself—Appendix 2

IEC 61508 is a generic standard which deals with the above. It can be used on its own or as a basis for developing industry sector specific standards (Chapters 8, 9 & 10). In attempting to fill the roles of being both a global template for the development of application specific standards, and being a standard in its own right, it necessarily leaves much to the discretion and interpretation of the user. IEC 61511 is a simplified form of IEC 61508 catering for the more consistent equipment architectures found in the process industries. This edition includes a new chapter on cybersecurity (Chapter 17). The topic has risen sharply in importance over the last few years and the requirement to address it in safety related studies is iterated in both IEC 61508 and also in IEC 61511.

One should bear in mind that the above documents are, largely, nonprescriptive guidance and a large amount of interpretation is required on the part of the user. There are few absolute right/wrong answers and, as always, the judgment of the professional (i.e., chartered) engineer must always prevail. In that respective, they might better be described as guidance documents rather than standards.

It is also vital to bear in mind that no amount of assessment will lead to enhanced integrity unless the assessment process is used as a tool during the design-cycle.

Now read on!

The 2010 Version of IEC 61508

The following is a brief summary of the main changes which brought about the 2010 version.

Architectural Constraints (Chapter 3)

An alternative route to the "safe failure fraction" (the so-called route 1_H) requirements was introduced (known as Route 2_H).

Route 2_H allows the "safe failure fraction" requirements to lapse providing that amount of redundancy (so called hardware fault tolerance) meets a minimum requirement and there is adequate user-based information providing failure rate data.

The meaning of "safe" failures in the formula for Safe Failure Fraction was emphasized as referring only to failures which force a "safe" state (e.g., spurious trip).

Security (Chapter 2)

Malevolent and unauthorized actions, as well as human error and equipment failure, can be involved in causing a hazard. They are to be taken account of, if relevant, in risk assessments.

Safety Specifications (Chapter 3)

There is more emphasis on the distinct safety requirements leading to separately defined design requirements.

Digital Communications (Chapter 3)

More detail in providing design and test requirements for "black box" and "white box" communications links.

ASICs and Integrated Circuits (Chapters 3 and 4)

More detailed techniques and measures are defined and described in Annexes to the Standard.

Safety Manual (Chapters 3 and 4)

Producers are required to provide a safety manual (applies to hardware and to software) with all the relevant safety-related information. Headings are described in Annexes to the Standard.

Synthesis of Elements (Chapter 3)

In respect of systematic failures, the ability to claim an increment of one SIL for parallel elements.

Software Properties of Techniques (Chapter 4)

New guidance on justifying the properties which proposed alternative software techniques should achieve in order to be acceptable.

Element (Appendix 8)

The introduction of a new term “element” (similar to a subsystem).

The 2016 Version of IEC 61511

The following is a brief summary of the main changes which have brought about the 2016 update.

The Safety Manual (IEC 65108 2010) is emphasized.

Procedures for competence are called for.

It is possible to claim up to one risk reduction layer within the process control system for the same hazard event when it is also the initiating event and two risk reduction layers if it is not part of the initiating cause (see Chapter 8).

The Architectures (i.e., Safe Failure Fraction) table is revised (see Chapter 8).

Acknowledgments

The authors would like to thank all the staff of ESC Ltd. for suggestions and support and, in particular, Simon Burwood, Dr Fan Ye, and Dr Hui Peng Li for their valuable contributions.

Thanks, also, to

Mr Colin Easton, of Prosalus, for extensive and useful inputs and guidance on cyber integrity.

Dr Tony Foord for constructive comments on Chapters 3 and 4 and for help with the original Chapter 14.

Mr Paul Reeve for comments on Chapter 7.

Mr Stephen Waldron, of JCB, and Mr Peter Stanton, of Railtrack, for help with Chapter 10.

Mike Dodson, Independent Consultant, of Solihull, for extensive comments and suggestions and for a thorough reading of the earlier manuscripts.

The authors are also grateful to Mirek Generowicz, Principal Consultant, I&E Systems Pty Ltd., Australia, for some useful comments on the 4th edition.

PART A

The Concept of Safety Integrity

In the first chapter we will introduce the concept of functional safety and the need to express targets by means of safety integrity levels. Functional safety will be placed in context, along with risk assessment, likelihood of fatality, and the cost of conformance.

The life-cycle approach, together with the basic outline of IEC 61508 (known as BS EN 61508 in the UK), will be explained.

CHAPTER 1

The Meaning and Context of Safety Integrity Targets

1.1 Risk and the Need for Safety Targets

There is no such thing as zero risk. This is because no physical item has zero failure rate, no human being makes zero errors, and no piece of software design can foresee every operational possibility.

Nevertheless public perception of risk, particularly in the aftermath of a major incident, often calls for the zero risk ideal. However, in general, most people understand that this is not practicable, as can be seen from the following examples of everyday risk of death from various causes:

All causes (mid-life including medical)	1×10^{-3} pa
All accidents (per individual)	5×10^{-4} pa
Accident in the home	4×10^{-4} pa
Road traffic accident	6×10^{-5} pa
Natural disasters (per individual)	2×10^{-6} pa

Therefore the concept of defining and accepting a tolerable risk for any particular activity prevails.

The actual degree of risk considered to be tolerable will vary according to a number of factors such as the degree of control one has over the circumstances, the voluntary or involuntary nature of the risk, the number of persons at risk in any one incident, and so on. This partly explains why the home remains one of the highest areas of risk to the individual in everyday life since it is there that we have control over what we choose to do and are therefore prepared to tolerate the risks involved.

A safety technology has grown up around the need to set target risk levels and to evaluate whether proposed designs meet these targets, be they process plant, transport systems, medical equipment, or any other application.

The Safety Critical Systems Handbook. https://doi.org/10.1016/B978-0-12-820258-6.00001-2

In the early 1970s people in the process industries became aware that, with larger plants involving higher inventories of hazardous material, the practice of learning by mistakes (if indeed we do) was no longer acceptable. Methods were developed for identifying hazards and for quantifying the consequences of failures. They were evolved largely to assist in the decision-making process when developing or modifying a plant. External pressures to identify and quantify risk were to come later.

By the mid 1970s there was already concern over the lack of formal controls for regulating those activities which could lead to incidents having a major impact on the health and safety of the general public. The Flixborough incident in June 1974, which resulted in 28 deaths, focused UK public and media attention on this area of technology. Many further events, such as that at Seveso (Italy) in 1976 through to the Piper Alpha offshore disaster and more recent Paddington (and other) rail incidents, have kept that interest alive and have given rise to the publication of guidance and also to legislation in the UK.

The techniques for quantifying the predicted frequency of failures are just the same as those previously applied to plant availability, where the cost of equipment failure was the prime concern. The tendency in the last few years has been towards a more rigorous application of these techniques (together with third-party verification) in the field of hazard assessment. They include Fault Tree Analysis, Failure Mode Effect Analysis, Common Cause Failure Assessment, and so on. These will be explained in Chapters 5 and 6.

Hazard assessment of process plant, and of other industrial activities, was common in the 1980s, but formal guidance and standards were rare and somewhat fragmented. Only Section 6 of the Health and Safety at Work Act 1974 underpinned the need to do all that is reasonably practicable to ensure safety. However, following the Flixborough disaster, a series of moves (including the Seveso directive) led to the CIMAH (Control of Industrial Major Accident Hazards) regulations, 1984, and their revised COMAH form (Control of Major Accident Hazards) in 1999. Appendix 9 provides an overview of this area and an outline of the contents needed in a COMAH report. The adoption of the Machinery Directive by the EU, in 1989, brought the requirement for a documented risk analysis in support of CE marking.

Nevertheless, these laws and requirements neither specify how one should go about establishing a target tolerable risk for an activity, nor address the methods of assessment of proposed designs, nor provide requirements for specific safety-related features within design.

The need for more formal guidance has long been acknowledged. Until the mid 1980s risk assessment techniques tended to concentrate on quantifying the frequency and magnitude of consequences arising from given risks. These were sometimes compared with loosely defined target values but, being a controversial topic, such targets (usually in the form of fatality rates) were not readily owned up to or published.

EN 1050 (Principles of Risk Assessment), in 1996, covered the processes involved in risk assessment but gave little advice on risk reduction. For machinery control EN 954-1

(see Chapter 10) provided some guidance on how to reduce risks associated with control systems but did not specifically include PLCs (programmable logic controllers) which were separately addressed by other IEC (International Electrotechnical Commission) and CENELEC (European Committee for Standardization) documents.

The proliferation of software during the 1980s, particularly in real time control and safety systems, focused attention on the need to address systematic failures since they could not necessarily be quantified. In other words while hardware failure rates were seen as a credibly predictable measure of reliability, software failure rates were generally agreed not to be predictable. It became generally accepted that it was necessary to consider qualitative defenses against systematic failures as an additional, and separate, activity to the task of predicting the probability of so-called random hardware failures.

In 1989, the HSE (Health and Safety Executive) published guidance which encouraged this dual approach of assuring functional safety of programmable equipment. This led to IEC work, during the 1990s, which culminated in the international safety Standard IEC 61508—the main subject of this book. The IEC Standard is concerned with electrical, electronic, and programmable safety-related systems where failure will affect people or the environment. It has a voluntary, rather than legal, status in the UK but it has to be said that to ignore it might now be seen as "not doing all that is reasonably practicable" in the sense of the Health and Safety at Work Act and a failure to show "due diligence." As use of the Standard becomes more and more widespread it can be argued that it is more and more "practicable" to use it. The Standard was revised and re-issued in 2010. Figure 1.1 shows how IEC 61508 relates to some of the current legislation.

The purpose of this book is to explain, in as concise a way as possible, the requirements of IEC 61508 and the other industry-related documents (some of which are referred to as second tier guidance) which translate the requirements into specific application areas.

The Standard, as with most such documents, has considerable overlap, repetition, and some degree of ambiguity, which places the onus on the user to make interpretations of the guidance and, in the end, apply his/her own judgment.

The question frequently arises as to what is to be classified as safety-related equipment. The term "safety-related" applies to any hard-wired or programmable system where a failure, singly or in combination with other failures/errors, could lead to death, injury, or environmental damage. The terms "safety-related" and "safety-critical" are often used and the distinction has become blurred. "Safety-critical" has tended to be used where failure alone, of the equipment in question, leads to a fatality or increase in risk to exposed people. "Safety-related" has a wider context in that it includes equipment in which a single failure is not necessarily critical whereas coincident failure of some other item leads to the hazardous consequences.

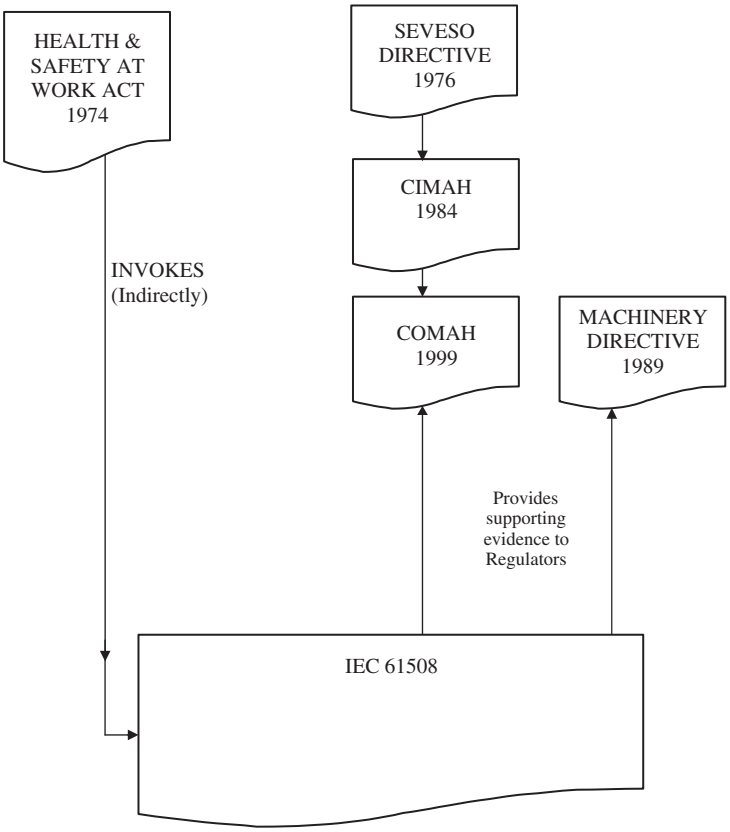

Figure 1.1: How IEC 61508 relates to some of the current legislation.

A piece of equipment, or software, cannot be excluded from this safety-related category merely by identifying that there are alternative means of protection. This would be to prejudge the issue and a formal safety integrity assessment would still be required to determine whether the overall degree of protection is adequate.

1.2 Quantitative and Qualitative Safety Target

In an earlier paragraph we introduced the idea of needing to address safety-integrity targets both quantitatively and qualitatively:

Quantitatively: where we predict the frequency of hardware failures and compare them with some tolerable risk target. If the target is not satisfied then the design is adapted (e.g., provision of more redundancy) until the target is met.

Qualitatively: where we attempt to minimize the occurrence of systematic failures (e.g., software errors) by applying a variety of defenses and design disciplines appropriate to the severity of the tolerable risk target.

It is important to understand why this twofold approach is needed. Prior to the 1980s, system failures could usually be identified as specific component failures (e.g., relay open circuit, capacitor short circuit, motor fails to start). However, since then the growth of complexity (including software) has led to system failures of a more subtle nature whose cause may not be attributable to a catastrophic component failure. Hence we talk of:

Random hardware failures: which are attributable to specific component failures and to which we attribute failure rates. The concept of "repeatability" allows us to model proposed systems by means of associating past failure rates of like components together to predict the performance of the design in question.

Systematic failures: which are not attributable to specific component failures and are therefore unique to a given system and its environment. They include design tolerance/timing-related problems, failures due to inadequately assessed modifications and, of course, software. Failure rates cannot be ascribed to these incidents since they do not enable us to predict the performance of future designs.

Quantified targets can therefore be set for the former (random hardware failures) but not for the latter. Hence the concept emerges of an arbitrary number of levels of rigor/excellence in the control of the design and operations. The ISO 9001 concept of a qualitative set of controls is somewhat similar and is a form of single "SIL." In the Functional Safety profession the practice has been to establish four such levels of rigor according to the severity of the original risk target.

During the 1990s this concept of safety-integrity levels (known as SILs) evolved and is used in the majority of documents in this area. The concept is to divide the "spectrum" of integrity into four discrete levels and then to lay down requirements for each level. Clearly, the higher the SIL then the more stringent the requirements become. In IEC 61508 (and in most other documents) the four levels are defined as shown in Table 1.1.

Table 1.1: Safety integrity levels (SILs).

Safety integrity level	Continuous and high demand rate (dangerous failures/hr)	Low demand rate (probability of failure on demand)
4	$\geq 10^{-9}$ to $< 10^{-8}$	$\geq 10^{-5}$ to $< 10^{-4}$
3	$\geq 10^{-8}$ to $< 10^{-7}$	$\geq 10^{-4}$ to $< 10^{-3}$
2	$\geq 10^{-7}$ to $< 10^{-6}$	$\geq 10^{-3}$ to $< 10^{-2}$
1	$\geq 10^{-6}$ to $< 10^{-5}$	$\geq 10^{-2}$ to $< 10^{-1}$

Note that had the high-demand SIL bands been expressed as "per annum," then the tables would appear numerically similar. However, being different parameters, they are NOT even the same dimensionally. Thus the "per hour" units are used to minimize confusion.

The reason for there being effectively two tables (high and low demand) is that there are two ways in which the integrity target may need to be described. The difference can best be understood by way of examples.

Consider the motor car brakes. It is the rate of failure which is of concern because there is a high probability of suffering the hazard immediately after each failure occurs. Hence we have the middle column of Table 1.1.

On the other hand, consider the motor car air bag. This is a low-demand protection system in the sense that demands on it are infrequent (years or tens of years apart). Failure rate alone is of little use to describe its integrity since the hazard is not incurred immediately after each failure occurs and we therefore have to take into consideration the test interval. In other words, since the demand is infrequent, failures may well be dormant and persist during the test interval. What is of interest is the combination of failure rate and down time and we therefore specify the probability of failure on demand (PFD): hence the right hand column of Table 1.1.

> *In IEC 61508 (clause 3.5.14 of part 4) the high demand definition is called for when the demand on a safety related function is greater than once per annum and the low demand definition when it is less frequent. However there is some debate on this issue and the authors believe that low demand might realistically be claimed when the demand rate is much less than the test frequency (typically an order of magnitude).*

In Chapter 2 we will explain the ways of establishing a target SIL and it will be seen that the IEC 61508 Standard then goes on to tackle the two areas of meeting the quantifiable target and addressing the qualitative requirements separately.

A frequent misunderstanding is to assume that if the qualitative requirements of a particular SIL are observed the numerical failure targets, given in Table 1.1, will automatically be achieved. This is most certainly not the case since the two issues are quite separate. The quantitative targets refer to random hardware failures and are dealt within Chapters 5 and 6. The qualitative requirements refer to quite different types of failure whose frequency is NOT quantified and are thus dealt with separately. The assumption, coarse as it is, is that by spreading the rigor of requirements across the range SIL 1 to SIL 4, which in turn covers the credible range of achievable integrity, the achieved integrity is likely to coincide with the measures applied.

A question sometimes asked is: If the quantitative target is met by the predicted random hardware failure probability then what allocation should there be for the systematic (software) failures? The target is to be applied equally to random hardware failures and to systematic failures. In other words the numerical target is not divided between the two but applied to the random hardware failures. The corresponding SIL requirements are then applied to the systematic failures. In any case, having regard to the accuracy of

quantitative predictions (see Chapter 6), the point may not be that important. The 2010 version implies this in 7.4.5.1 of Part 2.

The following should be kept in mind:

SIL 1: is relatively easy to achieve especially if ISO 9001 practices apply throughout the design providing that Functional Safety Capability is demonstrated.
SIL 2: is not dramatically harder than SIL 1 to achieve although clearly involving more review and test and hence more cost. Again, if ISO 9001 practices apply throughout the design, it should not be difficult to achieve.

(SILs 1 and 2 are not dramatically different in terms of the life-cycle activities)

SIL 3: involves a significantly more substantial increment of effort and competence than is the case from SIL 1 to SIL 2. Specific examples are the need to revalidate the system following design changes and the increased need for training of operators. Cost and time will be a significant factor and the choice of vendors will be more limited by lack of ability to provide SIL 3 designs.
SIL 4: involves state-of-the-art practices including "formal methods" in design. Cost will be extremely high and competence in all the techniques required is not easy to find. There is a considerable body of opinion that SIL 4 should be avoided and that additional levels of protection should be preferred.

It is reasonable to say that the main difference between the SILs is the quantification of random hardware failures and the application of the Safe Failure Fraction rules (see Chapter 3). The qualitative requirements for SILs 1 and 2 are very similar, as are those for SILs 3 and 4. The major difference is in the increment of rigor between SIL 2 and SIL 3.

Note, also, that as one moves up the SILs the statistical implications of verification become more onerous whereas the assessment becomes more subjective due to the limitations of the data available for the demonstration.

1.3 The Life-Cycle Approach

Section 7.1 of Part 1

The various life-cycle activities and defenses against systematic failures, necessary to achieve functional safety, occur at different stages in the design and operating life of equipment. Therefore it is considered a good idea to define (that is to say describe) a life cycle.

IEC 61508 is based on a safety life-cycle approach, describes such a model, and identifies activities and requirements based on it. It is important to understand this because a very

large proportion of safety assessment work has been (and often still is) confined to assessing whether the proposed design configuration (architecture) meets the target failure probabilities (dealt with later in Chapters 5 and 6 of this book). Because of systematic failures, modern guidance (especially IEC 61508) requires a much wider approach involving control over all of the life-cycle activities that influence safety integrity.

Figure 1.2 shows a simple life cycle very similar to the one shown in the Standard. It has been simplified for the purposes of this book.

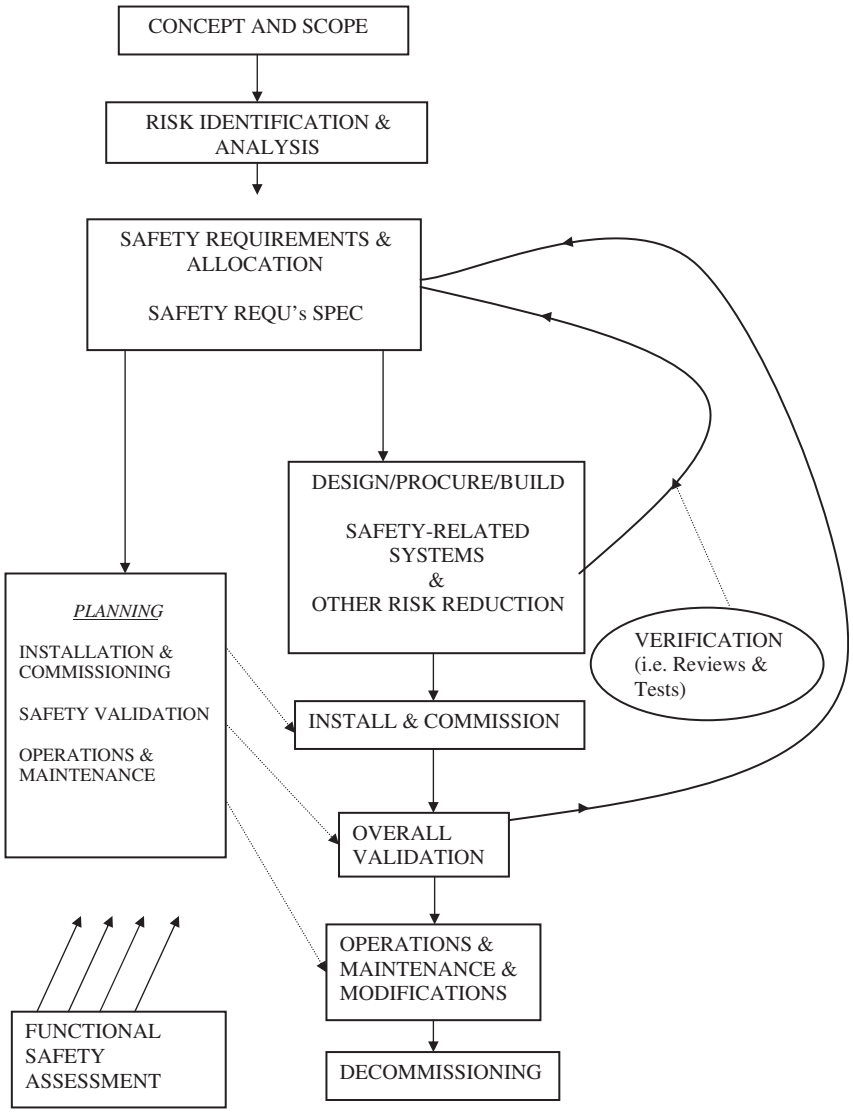

Figure 1.2: Safety life cycle.

As far as IEC 61508 is concerned this life cycle applies to all electrical and programmable aspects of the safety-related equipment. Therefore if a safety-related system contains an E/PE element then the Standard applies to all the elements of system, including mechanical and pneumatic equipment. There is no reason, however, why it should not also be used in respect of "other technologies" where they are used to provide risk reduction. For that reason the Gas Industry document IGEM/SR/15 is entitled "Integrity of safety-related systems in the gas industry" in order to include all technologies.

The IEC 61508 headings are summarized in the following pages and also map to the headings in Chapters 3 and 4. This is because the Standard repeats the process for systems hardware (Part 2) and for software (Part 3). IEC 65108 Part 1 lists these in its "Table 1" with associated paragraphs of text. The following text refers to the items in IEC 61508 Part 1 Table 1 and provides the associated paragraph numbers.

Concept and scope [Part 1—7.2 and 7.3]

Defines exactly what is the EUC (equipment under control) and the part(s) being controlled. Understands the EUC boundary and its safety requirements. Scopes the extent of the hazard and identification techniques (e.g., HAZOP). Requires a safety plan for all the life-cycle activities.

Hazard and risk analysis [Part 1—7.4]

This involves the quantified risk assessment by considering the consequences of failure (often referred to as HAZAN).

Safety requirements and allocation [Part 1—7.5 and 7.6]

Here we address the WHOLE SYSTEM and set maximum tolerable risk targets and allocate failure rate targets to the various failure modes across the system. Effectively this defines what the safety function is by establishing what failures are protected against and how. Thus the safety functions are defined and EACH has its own SIL (see Chapter 2).

Plan operations and maintenance [Part 1—7.7]

What happens in operations, and during maintenance, can effect functional safety and therefore this has to be planned. The effect of human error is important here as will be covered in Chapter 5. This also involves recording actual safety-related demands on systems as well as failures.

Plan the validation [Part 1—7.8]

Here we plan the overall validation of all the functions. It involves pulling together the evidence from the all the verification (i.e., review and test) activities into a coherent demonstration of conformance to the safety-related requirements.

Plan installation and commissioning [Part 1—7.9]

What happens through installation and commissioning can affect functional safety and therefore this has to be planned. The effect of human error is important here as will be shown in Chapter 5.

The safety requirements specification [Part 1—7.10]

Describes all the safety functions in detail.

ESC's SILComp® software generates a safety-related specification automatically based on data from SIL Targeting and Assessment (Verification). This is of particular use when managing a large number of SIFs.

Design and build the system [Part 1—7.11 and 7.12]

This is called "realization" in IEC 61508. It means creating the actual safety systems be they electrical, electronic, pneumatic, and/or other failure protection levels (e.g., physical bunds or barriers).

Install and commission [Part 1—7.13]

Implement the installation and create records of events during installation and commissioning, especially failures.

Validate that the safety-systems meet the requirements [Part 1—7.14]

This involves checking that all the allocated targets (above) have been met. This will involve a mixture of predictions, reviews, and test results. There will have been a validation plan (see above) and there will need to be records that all the tests have been carried out and recorded for both hardware and software to see that they meet the requirements of the target SIL. It is important that the system is re-validated from time to time during its life, based on recorded data.

Operate, maintain, and repair [Part 1—7.15]

Clearly operations and maintenance (already planned above) are important. Documentation, particularly of failures, is important.

Control modifications [Part 1—7.16]

It is also important not to forget that modifications are, in effect, re-design and that the life-cycle activities should be activated as appropriate when changes are made.

Disposal [Part 1—7.17]

Finally, decommissioning carries its own safety hazards which should be taken into account.

Verification [Part 1—7.18]

Demonstrating that all life-cycle stage deliverables were met in use.

Functional safety assessments [Part 1—8]

Carry out assessments to demonstrate compliance with the target SILs (see Section 2.3 of this book for the extent of independence according to consequences and SIL).

1.4 Steps in the Assessment Process

The following steps are part of the safety life-cycle (functional safety assessment).

Step 1. Establish Functional Safety Capability (i.e., Management)

Whereas Steps 2—7 refer to the assessment of a system or product, there is the requirement to establish the FUNCTIONAL SAFETY CAPABILITY of the assessor and/or the design organization. This is dealt with in Section 2.3 and by means of Appendix 1.

Step 2. Establish a Risk Target

ESTABLISH THE RISK TO BE ADDRESSED by means of techniques such as formal hazard identification or HAZOP whereby failures and deviations within a process (or equipment) are studied to assess outcomes. From this process one or more hazardous events may be revealed which will lead to death or serous injury.

SET MAXIMUM TOLERABLE FAILURE RATES by carrying out a quantified risk assessment based on a maximum tolerable probability of death or injury, arising from the event in question. This is dealt with in the next Chapter and takes into account how many simultaneous risks to which one is exposed in the same place, the number of fatalities and so on.

Step 3. Identify the Safety Related Function(s)

For each hazardous event it is necessary to understand what failure modes will lead to it. In this way the various elements of protection (e.g. control valve AND relief valve AND slamshut valve) can be identified. The safety protection system for which a SIL is needed can then be identified.

Step 4. Establish SILs for the Safety-Related Elements

Both the NUMERICAL ASSESSMENT, LOPA and RISK GRAPH methods are described in Chapter 2 and examples are given in Chapter 13.

Step 5. Quantitative Assessment of the Safety-Related System

Reliability modeling is needed to assess the failure rate or probability of failure on demand of the safety-related element or elements in question. This can then be compared with the target set in Step 3. Chapters 5 and 6 cover the main techniques.

Step 6. Qualitative Assessment Against the Target SILs

The various requirements for limiting systematic failures are more onerous as the SIL increases. These cover many of the life-cycle activities and are covered in Chapters 3 and 4.

Step 7. Establish ALARP

It is not sufficient to establish, in Step 4, that the quantitative failure rate (or the PFD) has been met. Design improvements which reduce the failure rate (until the Broadly Acceptable failure rate is met) should be considered and an assessment made as to whether these are "as low as reasonably practicable". This is covered in Section 2.2.

It is worth noting, at this point, that conformance to a SIL requires that all the Steps are met. If the quantitative assessment (Step 5) indicates a given SIL then this can only be claimed if the qualitative requirements (Step 6) are also met.

Part 1 clause 8 of IEC 61508 (Functional Safety Assessment) addresses this area. FSA should be done at all lifecycle phases (not just Phase 9, Realization). There are minimum levels of independence of the assessment team from the system/company being assessed, depending on the SIL involved. In summary these are:

SIL	Consequence	Assessed by
4	Many deaths**	Independent organisation
3*	More than one death**	Independent department
2*	Severe injury or one death	Independent person
1	Minor injury	Independent person

*Add one level if there is lack of experience, unusual complexity or novel design.
**Not quantified in the standard.

Typical headings in an assessment report would be:

- Hazard scenarios and associated failure modes
- SIL targeting
- Random hardware failures
- ALARP
- Architectures (SFF)
- Life-cycle activities
- Functional safety capability
- Recommendations.

1.5 Costs

The following questions are often asked:

"What is the cost of applying IEC 61508?"
"What are the potential savings arising from its use?"
"What are the potential penalty costs of ignoring it?"

1.5.1 Costs of Applying the Standard

Although costs will vary considerably, according to the scale and complexity of the system or project, the following typical resources have been expended in meeting various aspects of IEC 61508.

Full Functional Safety Capability (now called Functional Safety Management) including implementation on a project or product—30 to 60 man-days + several £'000 for certification by an accredited body (i.e. SIRA).
Product or Project Conformance (to the level of third-party independent assessment)—10 to 20 man-days + a few £'000 consultancy.

Elements within this can be identified as follows:

Typical SIL targeting with random hardware failures assessment and ALARP—2 to 6 man-days.
Assessing the safe failure fraction of an instrument (one or two failure modes)—1 to 3 man-days.
Bringing an ISO 9001 management system up to IEC61508 functional safety capability—5 man-days for the purpose of a product demonstration where evidence of only random hardware failures and safe failure fraction are being offered, 20 to 50 man-days for the purpose of an accredited Functional Safety Capability certificate.

1.5.2 Savings from Implementing the Standard

For some time there has an intangible but definite benefit due to enhanced credibility in the market place. Additional sales vis à vis those who have not demonstrated conformance are likely. However, the majority of instrument and system providers now see it as necessary to demonstrate conformance to some SIL and thus it becomes a positive disadvantage not to do so.

Major savings are purported to be in reduced maintenance for those (often the majority) systems which are given low SIL targets. This also has the effect of focusing the effort on the systems with higher SIL targets.

1.5.3 Penalty Costs from Not Implementing the Standard

The manufacturer and the user will be involved in far higher costs of retrospective redesign if subsequent changes are needed to meet the maximum tolerable risk.

The user could face enormous legal costs in the event of a major incident which invokes the H&SW Act especially if the Standard had not been applied when it was reasonably practicable to have done so.

1.6 The Seven Parts of IEC 61508

Now that we have introduced the concept of safety integrity levels and described the life-cycle approach it is now appropriate to describe the structure of the IEC 61508 Standard. Parts 1−3 are the main parts (Figure 1.3) and Parts 4−7 provide supplementary material.

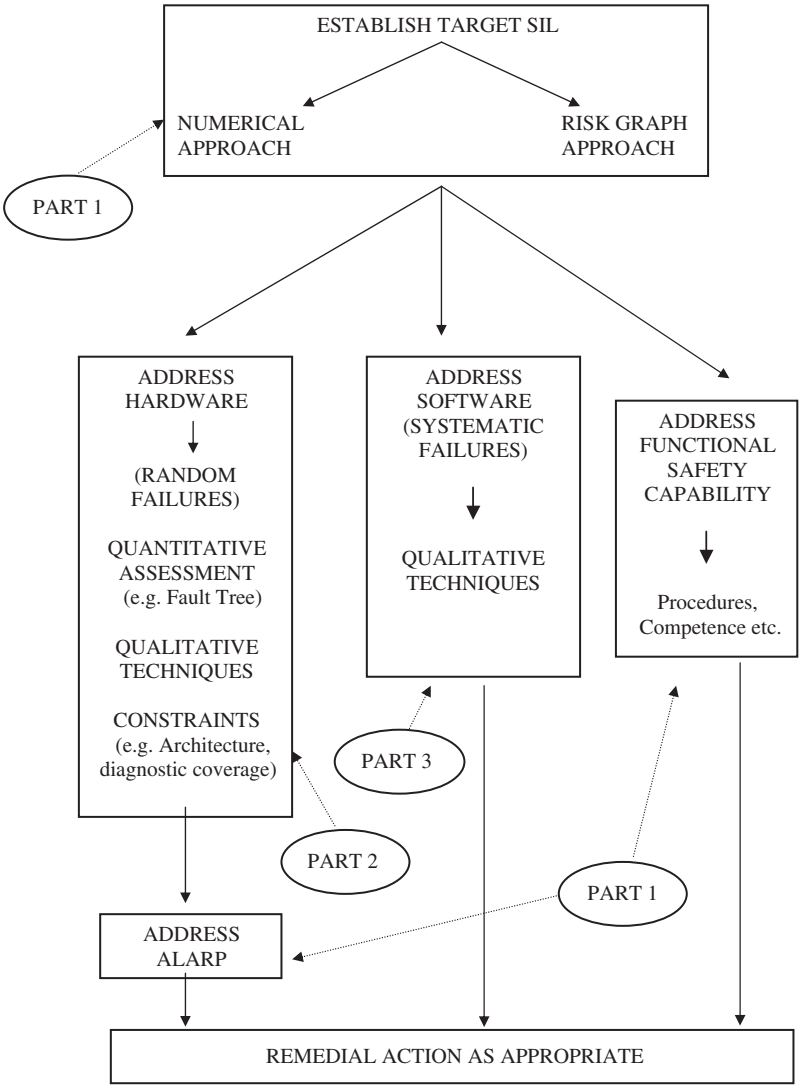

Figure 1.3: The parts of the standard.

The general strategy is to establish SIL targets, from hazard and risk analysis activities, and then to design the safety-related equipment to an appropriate integrity level taking into account random and systematic failures and also human error.

Examples of safety-related equipment might include:

Shutdown systems for processes
Interlocks for dangerous machinery
Fire and gas detection
Instrumentation
Programmable controllers
Railway signaling
Boiler and burner controls
Industrial machinery
Avionic systems
Leisure items (e.g. fairground rides)
Medical equipment (e.g. oncology systems).

Part 1 is called "General Requirements." In actual fact it covers:

(i) General functional safety management, dealt with in Chapter 2 and Appendix 1 of this book. This is the management system (possibly described in one's quality management system) which lays down the activities, procedures and skills necessary to carry out the business of risk assessment and of designing to meet integrity levels.
(ii) The life-cycle, explained above, and the requirements at each stage, are central to the theme of achieving functional safety. It will dominate the structure of several of the following Chapters and Appendices.
(iii) The definition of SILs and the need for a hazard analysis in order to define an SIL target.
(iv) The need for competency criteria for people engaged in safety-related work, also dealt with in Chapter 2 of this book.
(v) Levels of independence of those carrying out the assessment. The higher the SIL the more independent should be the assessment.
(vi) There is an Annex in Part 1 (informative only) providing a sample document structure for a safety-related design project.

Part 2 is called "Requirements for E/E/PES safety-related systems." What this actually means is that Part 2 is concerned with the hardware, rather than the software, aspects of the safety-related system. It covers:

(i) The life-cycle activities associated with the design and realization of the equipment including defining safety requirements, planning the design, validation, verification, observing architectural constraints, fault tolerance, test, subsequent modification (which will be dealt with in Chapter 3).

(ii) The need to assess (i.e. predict) the quantitative reliability (vis à vis random hardware failures) against the SIL targets in Table 1.1. This is the reliability prediction part of the process and is covered in Chapters 5 and 6.
(iii) The techniques and procedures for defending against systematic hardware failures.
(iv) Architectural constraints vis à vis the amount of redundancy applicable to each SIL. Hence, even if the above reliability prediction indicates that the SIL is met, there will still be minimum levels of redundancy. This could be argued as being because the reliability prediction will only have addressed random hardware failures (in other words those present in the failure rate data) and there is still the need for minimum defenses to tackle the systematic failures.
(v) Some of the material is in the form of annexes.

Chapter 3 of this book is devoted to summarizing Part 2 of IEC 61508.

Part 3 is called "Software requirements." As the title suggests this addresses the activities and design techniques called for in the design of the software. It is therefore about systematic failures and no quantitative prediction is involved.

(i) Tables indicate the applicability and need for various techniques at each of the SILs.
(ii) Some of the material is in the form of annexes.

Chapter 4 of this book is devoted to summarizing Part 3 of IEC 61508.

Part 4 is called "Definitions and abbreviations". This book does not propose to offer yet another list of terms and abbreviations beyond the few terms in Appendix 8. In this book the terms are hopefully made clear as they are introduced.

Part 5 is called "Examples of methods for the determination of safety-integrity levels".

As mentioned above, the majority of Part 5 is in the form of seven Annexes which are informative rather than normative:

(i) Annex A covers the general concept of the need for risk reduction through to the allocation of safety requirements, which is covered in Chapter 2 of this book.
(ii) Annex B covers methods for determining safety integrity level targets.
(iii) Annex C covers the application of the ALARP (as low as reasonably practicable) principle, which is covered in Section 2.2 of this book.
(iv) Annex D covers the mechanics of quantitatively determining the SIL levels, which is covered in Section 2.1 of this book.
(v) Annex E covers a qualitative method (risk graph) of establishing the SIL levels, which is also covered in Chapter 2 of this book.
(vi) Annex F covers Semi-quantitative LOPA (Chapter 2 of this book).
(vii) Annex G describes an alternative qualitative method, "Hazardous event severity matrix".

Part 6 is called "Guidelines on the application of Part 2 and 3". This consists largely of informative annexes which provide material on:

(i) Calculating hardware failure probabilities (low and high demand).
(ii) Common cause failure, which is covered in Chapter 5 of this book.
(iii) Diagnostic coverage, which is covered in Chapter 3 of this book.
(iv) Applying the software requirements tables (of Part 3) for SILs 2 and 3, which is covered in Chapter 4 of this book.

As mentioned above, the majority of Part 6 is in the form of Annexes which are informative rather than normative.

Part 7 is called "Overview of techniques and measures". This is a reference guide to techniques and measures and is cross-referenced from other parts of the Standard. This book does not repeat that list but attempts to explain the essentials as it goes along.

The basic requirements are summarized in Figure 1.4.

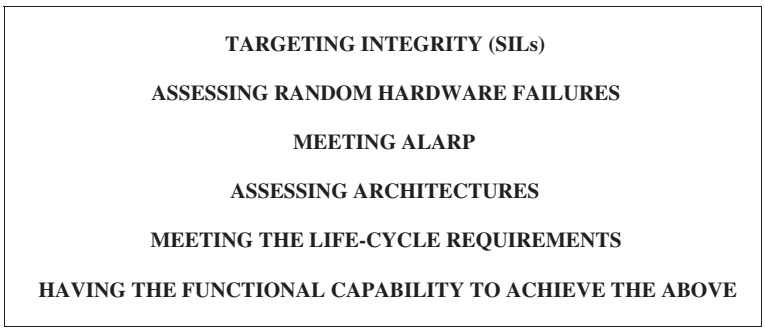

Figure 1.4: Summary of the requirements.

1.7 HAZOP (Hazard and Operability Study)

In the process industry it is normal to undertake a HAZOP to identify potential hazards that may require mitigation (e.g., instrumented protection systems) in order to reduce risk to persons, environment, or plant.

This book concerns functional safety assessment and Figure 1.5 illustrates how HAZOP and HAZID (Hazard Identification) studies provide the "trigger" for the functional safety assessment of plant and items of mitigation (ie safety related systems).

HAZOP is a study of the hazards associated with the actual process equipment in a plant whereas HAZID is a wider study which embraces all hazards associated with the whole plant (e.g., chemical, process, adjacent plant etc.). This is addressed in BS EN 61882 2016 Hazard and Operability Studies (HAZOP).

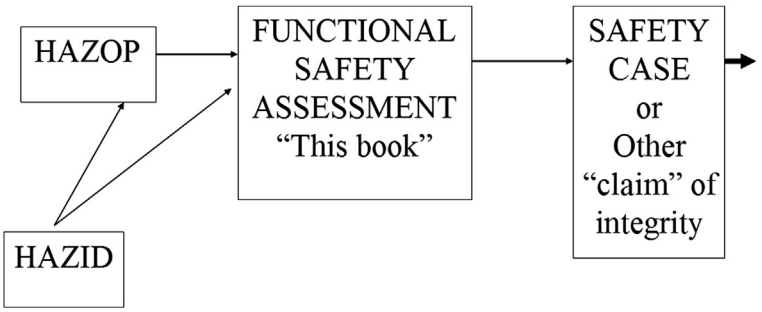

Figure 1.5: The context of HAZOP.

1.7.1 Objectives of a HAZOP

A HAZOP is a structured brainstorming meeting, set up to identify potential undesirable events that may create hazards or operability problems, i.e. risks to personal safety and potential damage to assets, the environment and the reputation. The identified hazards, along with any actions for further investigation, and other relevant supporting information are recorded on HAZOP worksheets.

A HAZOP is generally required to be carried out at various points during the design of the facility. It is anticipated that a HAZOP carried out at the early stages of design will typically produce numerous actions. As the design progresses to the later stages, it is anticipated the actions should begin to be closed out and prior to start-up there should be no outstanding actions.

The HAZOP methodology process is described in IEC 61882. The basic approach is to divide the plant process into convenient nodes then to go through each node using guide words to identify causes that could produce an undesirable event. These are recorded along with any intended safe guards and actions recorded if it is thought that additional safe guards may be required.

1.7.2 HAZOP Study Team

It is important that the proposed HAZOP team is made up of personnel who will bring the best balance of knowledge and experience, of the type of plant being considered, to the study. Table 1.2 outlines the typical minimum requirements of the proposed workshop team and corresponding roles and responsibilities.

Process experts of other disciplines should be available to be called upon as required:

- Rotating Machinery
- Maintenance
- Corrosion and materials
- Mechanical

Table 1.2: HAZOP team typical roles and responsibilities.

Responsibility	Role
To facilitate a robust consensus-based decision-making process within the HAZOP team Ensuring the HAZOP workshop is performed in accordance with the agreed Terms of Reference Ensuring the HAZOP workshop is accurately documented *Note: The chairmen would not normally be an expert for the process in question*	HAZOP chair
Documenting the proceedings of the HAZOP workshop in an efficient and accurate manner under the guidance of the chair Producing and maintaining an accurate record of the HAZOP workshop	HAZOP scribe
To provide expertise, experience and understanding of:	
The design of the plant under analysis including equipment; design limits; materials of construction and condition of equipment	Typically the project or site process engineer
The plant controls and instrumentation	Typically the project or site electrical, controls and instrumentation (EC&I) engineers
Day-to-day operation of the plant under analysis	Typically a senior operations representative
Process safety and the major hazards associated with plant under analysis	Typically the project or site process safety engineer

1.7.3 Typical Information Used in the HAZOP

The following documentation should generally be provided to each team member (as a minimum) for the workshop:

- Cause and Effects (C&E) Charts (or equivalent) for all SIFs (safety instrumented functions) under consideration;
- Accurate P&IDs, which reflect the as-built, or current design, status of the plant under analysis.
- The following documentation should be made available to the workshop team (but not necessarily to each team member individually):
  - Operating procedures;
  - Pressure Safety Valve (PSV) design data;
  - Vessel/piping design data;
  - Consequence analysis studies;
  - Plant layout drawings/Plot plans.

1.7.4 Typical HAZOP Worksheet Headings

The following sections detail each of the headings of the HAZOP worksheet.

Design Intent

Problems in the process normally arise from deviation outside the intended operating envelope. It is therefore important to record the design intent so that deviations can be identified from it. These deviations can be expressed in terms of temperature, level, pressure etc.

Nodes

The facility under analysis is divided into sections known as nodes with certain defined boundaries. Prior to the review of each node the design intent should be explained by the engineer(s). The size of each node is dependent on the batch process sequences, complexity of the system and any natural divisions in rating, pump suction, discharge piping and tanks/vessel/reactors etc.

Parameter/Guidewords

Each node is analyzed to determine the potential undesirable event(s) associated with that section of the facility. This is achieved by considering a given parameter and how that parameter can deviate (guidewords) from the design intent. A list of the parameters and typical guidewords used in the study which should be agreed by all parties (Table 1.3).

Table 1.3: Typical guidewords.

Parameter	Guideword
Flow	High, low, reverse
Temperature	High, low
Pressure	High, low
Level	High, low
Power	High, low
Mixing	High, low
Reaction	High, low
Composition	i.e., Wrong composition, wrong chemical, contamination
Contamination	—
Maintenance	Maintainability of equipment i.e. isolation capabilities etc.
Start-up/Shut-down	—

Causes

When the potential for an undesirable event exists, all causes of such an event are listed.

Consequence

The consequence of the undesirable event is stated. It is important to note that no credit can be given to any safeguards or means of mitigating this event at this stage.

Safeguards

The safeguards that can prevent the cause of the undesirable event occurring or mitigate the consequences are identified.

Action Required

Any actions associated with the section of facilities or undesirable event.

1.7.5 Risk Ranking

The priority assigned to a recommendation is generally defined by combining the severity and likelihood using a Risk Matrix.

The team can examine hazards with a safety, environmental or financial consequence. When no consequences are identified, no assessment is conducted as a result.

1.7.6 Quantifying Risk

It should be noted that this mythology was developed in 1970 when risks were generally dealt with in a qualitative approach therefore if there were a number of causes that could cause the same hazard each cause could be considered separately. Since 2000 due to IEC 61508 and IEC 61511 the risk should be quantified. Thus for any specific hazard all causes need to be considered. Unfortunately the majority of practitioners still use the existing approach of considering 'causes' in each node regardless of where the hazard occurs and regardless of other causes that could cause the same hazard, whereas a more suitable approach, with the latest requirements, would look for 'hazards' in each node then identify ALL causes inside and outside the node that can cause this hazard. This would make integrity assessment more straight forward.

The rest of this book deals with all aspects of functional safety assessment.

The ESC ProSET® software tool provides a convenient program for recording HAZOP and HAZID findings and can collate all causes associated with each specific hazard.

1.8 HAZIDS, CHAZOPs and SIMOPS

HAZID (Hazard Identification) is a qualitative technique for the early identification of potential hazards and threats affecting people, the environment, assets, or reputation. The process is identical to the HAZOP but whereas the HAZOP is solely considering the plant process equipment; the HAZID is considering all types of potential hazards that could affect the plant e.g., chemical spillage, dropped objects, etc.

The CHAZOP is a similar process to HAZOP but focuses on the potential hazards associated with the operations of the **control system** exclusive from the function of the underlying processes that they manipulate and maintain. The process focuses on whether the inputs/outputs of the control system are adequate for the process that they are applied to and, secondly, to evaluate whether the architect of the control system is properly integrated and does not have the potential to cause unacceptable plant operation or hazards. A CHAZOP provides a measure of the safety risks that are associated with a particular control system, and as such, it provides a measure of the control system's effectiveness.

The SIMOPS (Simultaneous Operations) procedure is similar to the HAZOP but is considering simultaneous operations during construction and commissioning to ensure sound safe work practices and procedures are applied during these simultaneous operations.

CHAPTER 2

Meeting IEC 61508 Part 1

Part 1 of the Standard addresses the need for:

- Setting integrity (SIL) targets
- The ALARP concept (by inference)
- Capability to design, operate, and maintain for functional safety
- Establishing competency
- Hierarchy of documents

The following sections summarize the main requirements:

2.1 Establishing Integrity Targets

Assessing quantified integrity targets is an essential part of the design process (including retrospective safety studies). This leads to:

- A quantified target against which one predicts the rate of random hardware failures and establishes ALARP (as low as reasonably practicable).
- A SIL band for mandating the appropriate rigor of life cycle activities.

The following paragraphs describe how a SIL target is established.

2.1.1 The Quantitative Approach

(a) Maximum Tolerable Risk

In order to set a quantified safety integrity target, a target Maximum Tolerable Risk is needed. It is therefore useful to be aware of the following rates:

All accidents (per individual)	5×10^{-4} pa
Natural disasters (per individual)	2×10^{-6} pa
Accident in the home	4×10^{-4} pa
Worst case Maximum Tolerable Risk in HSE R2P2 document	10^{-3} pa
"Very low risk" as described in HSE R2P2 document (i.e., boundary between Tolerable and Broadly Acceptable)	10^{-6} pa

The Safety Critical Systems Handbook. https://doi.org/10.1016/B978-0-12-820258-6.00002-4

"Individual Risk" is the frequency of fatality for a hypothetical person with respect to a specific hazard. This is different from "Societal Risk," which takes account of multiple fatalities. Society has a greater aversion to multiple fatalities than single ones in that killing 10 people in a single incident is perceived as worse than 10 separate single fatalities.

Table 2.1 shows the limits of tolerability for "Individual Risk" and is based on a review of HSE's "Reducing risk, protecting people, 2001 (R2P2)" and HSG87. The former indicates a Maximum Tolerable Risk to an employee of 10^{-3} per annum for all risks combined. The actual risk of accidents at work per annum is well below this. Generally, guidance documents recommend a target of 10^{-4} per annum for all process related risks combined, leaving a margin to allow for other types of risk.

At the lower end of the risk scale, a Broadly Acceptable Risk is nearly always defined. This is the risk below which one would not, normally, seek further risk reduction. It is approximately two or three orders of magnitude less than the total of random risks to which one is exposed in everyday life.

Table 2.1: Target individual risks.

	HSE R2P2	Generally used for functional safety
Maximum Tolerable Individual Risk (per annum)		
Employee	10^{-3}	10^{-4}
Public	10^{-4}	10^{-5}
Broadly Acceptable Risk (per annum)		
Employee and public	10^{-6}	10^{-6}

It is important to note that the Individual Risk and the Societal Risk calculations are fundamentally different. Thus the starting points for Maximum Tolerable Risk, in the case of a single fatality, do not immediately coincide, which will be elaborated in Section 2.4.

Scenarios, such as sites, usually imply a risk to the same (more or less) groups of individuals (be it on-site or off-site) at any time. "Distributed" risks, for example, pipelines across wide areas, rail journeys, tunnels, with rapidly changing identities of individuals are the scenarios for which the involuntary risk approach becomes limited. An individual may be exposed for 2 min per annum (traveling through a tunnel) whereas, at any moment, there may be 100 people at risk. The Societal Risk approach (Section 2.4) is then more appropriate.

There is a body of opinion that multiple fatalities should also affect the choice of maximum tolerable Individual Risk. The targets in Table 2.2 reflect an attempt to take account of these concerns in a relatively simple way by adjusting the Individual Risk targets from Table 2.1.

More complex calculations for Societal Risk (involving F–N curves) are sometimes addressed by specialists as are adjustments for particularly vulnerable sections of the community (disabled, children etc.).

Table 2.2: Target fatality risks.

	1–2 fatalities	3–5 fatalities	6 or more fatalities
Maximum Tolerable Individual Risk (per annum)			
Employee (Voluntary)	10^{-4}	3×10^{-5}	10^{-5}
Public (Involuntary)	10^{-5}	3×10^{-6}	10^{-6}
Broadly Acceptable Risk (per annum)			
Employee and public	10^{-6}	3×10^{-7}	10^{-7}

The location, that is, site or part of a site, for which a risk is being addressed, may be exposed to multiple potential sources of risk. The question arises as to how many potential separate hazards an individual (or group) in any one place and time is exposed to. Therefore, in the event of exposure to several hazards at one time, one should seek to allow for this by specifying a more stringent target for each hazard. For example, a study addressing a multirisk installation might need to take account of an order of magnitude of sources of risk. On the other hand, an assessment of a simple district pressure regulator valve for the local distribution of natural gas implies a limited number of sources of risk (perhaps only one).

A typical assessment confined to employees on a site might use the recommended 10^{-4} pa Maximum Tolerable Risk (for 1–2 fatalities) but may address 10 sources of risk to an individual in a particular place. Thus, an average of 10^{-5} pa would be used as the Maximum Tolerable Risk across the 10 hazards and, therefore, for each of the 10 safety functions involved. By the same token, the Broadly Acceptable Risk would be factored from 10^{-6} pa to 10^{-7} pa.

The question arises of how long an individual is exposed to a risk. Earlier practice has been to factor the maximum tolerable failure rate by the proportion of time it offers the risk (for example, an enclosure which is only visited 2 hrs per week). However, that approach would only be valid if persons (on-site) suffered no other risk outside that 2 hrs of his/her week. In case of off-site, the argument might be different in that persons may well only be at risk for a proportion of the time. Thus, for on-site personnel, the proportion of employee exposure time should be taken as the total working proportion of the week.

Despite the widely published figures for Maximum Tolerable Risk (e.g., Table 2.2), the UK HSE sometimes press for a Maximum Tolerable Risk to be targeted at a lower level nearer to the Broadly Acceptable level (e.g., an order of magnitude). This, however, is a controversial area. In the authors' opinion, whatever may the starting point be, the ALARP calculation will, in any case, cause the risk to be reduced to an appropriate level.

Table 2.3 caters for the lesser consequence of injury. Targets are set in the same manner and integrity assessment is carried out as for fatality. In general, rates an order of magnitude larger are used for the targets.

Table 2.3: Target individual risks for injury.

Maximum Tolerable Risk (per annum)	
Employee (Voluntary)	10^{-3}
Public (Involuntary)	10^{-4}
Broadly Acceptable Risk (per annum)	
Employee and public	10^{-5}

In any event, the final choice of Maximum Tolerable Risk (in any scenario) forms part of the "safety argument" put forward by a system user. There are no absolute rules but the foregoing provides an overview of current practice.

(b) Maximum tolerable failure rate

This involves factoring the Maximum Tolerable Risk according to totally external levels of protection and to factors which limit the propagation to fatality of the event. Table 2.4 gives examples of the elements which might be considered. These are not necessarily limited to the items described below and the analyst(s) must be open ended in identifying and assessing the factors involved.

The maximum tolerable failure rate is then targeted by taking the Maximum Tolerable Risk and factoring it according to the items assessed. Thus, for the examples given in Table 2.4 (assuming a 10^{-5} pa involuntary risk):

$$\begin{aligned}\text{Maximum Tolerable Failure Rate} &= 10^{-5}\ \text{pa}/(0.6 \times 0.2 \times 0.7 \times 0.25 \times 0.9 \times 0.25) \\ &= \mathbf{2.1 \times 10^{-3}\ pa}\end{aligned}$$

Table 2.4: Factors leading to the maximum tolerable failure rate.

Factor involving the propagation of the incident or describing an independent level of protection	Probability (example)	This column is used to record arguments, justifications, references etc. to support the probability used
The profile of time at risk	60%	Quantifying whether the scenario can develop. This may be <100% as for example if: • flow, temp, pressure etc. profiles are only sufficient at specific times, for the risk to apply. • the process is only in use for specific periods.
Unavailability of separate mitigation fails (i.e., another level of protection)	20%	Mitigation outside the scope of this study and not included in the subsequent modeling which assesses whether the system meets the risk target. Examples are: • α down stream temp, pressure etc. measurement leading to manual intervention. • a physical item of protection (for example, vessel; bund) not included in the study.
Probability of the scenario developing	70%	Examples are: • the vessel/line will succumb to the over-temp, over pressure etc. • the release has an impact on the passing vehicle.
Person(s) exposed (i.e., being at risk)	25%	Proportion of time during which some person or persons are close enough to be at risk should the event propagate. Since a person may be exposed to a range of risks during the working week, this factor should not be erroneously reduced to the proportion of time exposed to the risk in question. If that were repeated across the spectrum of risks then each would be assigned an artificially optimistic target. The working week is approximately 25% of the time and thus that is the factor which would be anticipated for an on-site risk. In the same way, an off-site risk may only apply to a given individual for a short time.
Probability of subsequent ignition	90%	Quantifying whether the released material ignites/ explodes.
Fatality ensues	25%	The likelihood that the event, having developed, actually leads to fatality.

Example

A gas release (e.g., a natural gas holder overfill) is judged to be a scenario leading to a single on-site fatality and three off-site fatalities. Both on and off site, person(s) are believed to be exposed to that one risk from the installation.

On site

Proportion of time system can offer the risk	75%	40 weeks pa
Probability of ignition	5%	Judgment
Person at risk	25%	Working week i.e., 42 hrs/168 hrs
Probability of fatality	75%	Judgment

From Table 2.2, the Maximum Tolerable Risk is 10^{-4} pa. Thus, the maximum tolerable failure rate (leading to the event) is calculated as:

$$10^{-4}\ \text{pa}/(0.75 \times 0.05 \times 0.25 \times 0.75) = 1.4 \times 10^{-2}\ \text{pa}$$

Off site

Proportion of time system can offer the risk	75%	40 weeks pa
Probability of ignition	5%	Judgment
Person(s) at risk	33%	Commercial premises adjoin
Probability of three fatalities	10%	Offices well protected by embankments

From Table 2.2 the Maximum Tolerable Risk is 3×10^{-6} pa. Thus the maximum tolerable failure rate (leading to the event) is calculated as:

$$3 \times 10^{-6}\ \text{pa}/(0.75 \times 0.05 \times 0.33 \times 0.1) = 2.4 \times 10^{-3}\ \text{pa}$$

Thus, 2.4×10^{-3} pa, being the more stringent of the two, is taken as the maximum tolerable failure rate target.

(c) Safety integrity levels (SILs)

Notice that only now is the SIL concept introduced. The foregoing is about risk targeting but the practice of jumping immediately to a SIL target is a dangerous approach.

Furthermore, it is necessary to understand why there is any need for a SIL concept when we have numerical risk targets against which to assess the design. If the assessment were to involve only traditional reliability prediction, wherein the predicted hardware reliability is compared with a target, there would be no need for the concept of discrete SILs. However, because the rigor of adherence to design/quality assurance activities cannot be quantified, a number of discrete levels of "rigor," which cover the credible range of