The Institution of Engineering and Technology

National Cyber Security Centre
a part of GCHQ

Code of Practice

Cyber Security and Safety

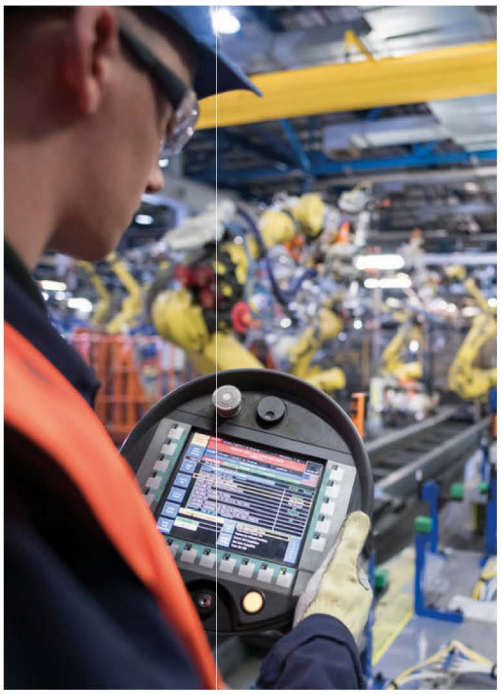

The Institution of Engineering and Technology

Code of Practice

Cyber Security and Safety

Publication Information

Published by The Institution of Engineering and Technology, London, United Kingdom
The Institution of Engineering and Technology is registered as a Charity in England & Wales (no. 211014) and Scotland (no. SCO38698).

The Institution of Engineering and Technology

The Institution of Engineering and Technology is the institution formed by the joining together of the IEE (The Institution of Electrical Engineers) and the IIE (The Institution of Incorporated Engineers).

First published 2020 (978-1-83953-318-1)

Copies of this publication may be obtained from:

PO Box 96
Stevenage
SG1 2SD, UK
Tel: +44 (0)1438 767328
Email: sales@theiet.org
https://electrical.theiet.org/guidance-codes-of-practice/

ISBN 978-1-83953-318-1 (paperback)
ISBN 978-1-83953-319-8 (electronic)

Typeset in the UK by the Institution of Engineering and Technology, Stevenage
Printed in the UK by Sterling Press

Contents

Lists of Figures and Tables

List of Figures

List of Tables

Acknowledgements

Background to the Code of Practice

This Code of Practice was developed following discussions within two of The Institution of Engineering and Technology's (IET's) Technical and Professional Networks (TPNs), one covering 'Functional Safety' and the other 'Cyber Security'. They identified a clear need to provide a Code of Practice that addresses the through-life design and management essentials for safety and cyber security and that would help promote and improve good practice. A workshop was undertaken involving a wide range of stakeholders, which identified the key principles.

The initial principles were revised and refined under the guidance of a technical committee by the authors to address the interactions and dependencies between the disciplines. It was decided that this Code should be international in outlook but contain specific UK regulatory flavours in the footnotes and additional guidance.

Further details on the activities of The IET and its TPNs can be found at https://www.theiet.org/.

It is the intention of The IET and NCSC to review this work regularly, with a view to keeping it up to date, relevant and valuable.

Acknowledgements

The IET would like to acknowledge the following people and organizations for their input into this Code of Practice:

Authors

Phil Williams, Engineer for Safety Limited, Lead author

Mike St. John-Green, Mike StJohn-Green Consulting Limited, Co-author/Technical Committee chair

IET Technical and Professional Networks

- Andy German Functional Safety TPN Chair and Technical committee member
- Richard Piggin Cyber security TPN Chair and Technical committee member

Technical Committee – organizations represented

- Atkins
- BAE Systems
- Cisco
- City University of London
- The Health and Safety Executive (HSE)
- Information Assurance Strategies Ltd.
- LDRA Limited
- Ministry of Defence
- Office for Nuclear Regulation (ONR)
- Rapiscan Systems
- Siemens
- University of Southampton
- University of Hertfordshire
- University of York

Foreword

Computerized systems are taking on an increasing role in performing vital safety-related functions, designed to protect human lives. Already, such systems are controlling the safe operation of industrial sites processing and storing dangerous chemicals, and play a key role in the safety of aviation and rail transportation, etc.

There is an increasing recognition that computerized safety systems could, potentially, be adversely affected by a cyber incident – either as a side-effect of a compromise not intended by the perpetrators to affect safety, or as a result of highly targeted cyber attack, specifically aimed at reducing the effectiveness of safety mechanisms. This is now more than just a possibility; in December 2017, for example, the NCSC became aware of malware dubbed TRITON, which was targeting the Triconex industrial safety controller, used in many installations worldwide.

Successful management of cyber-related risks to safety is based on the same fundamental principles that underpin effective cyber risk management more generally. However, an integrated approach is needed, which combines the established good practices of both the security and safety communities. There are recognized challenges associated with achieving such an integrated approach and the NCSC is keen to work with others to develop additional guidance. The NCSC has been pleased to support the IET in the production of this Code of Practice for cyber security and safety, by providing advice and input on the cyber security aspects of the document.

Carolyn A
NCSC Chief Engineer

National Cyber Security Centre (NCSC)

The NCSC was launched in 2016 in order to provide a single point of contact on cyber security matters for Small and Medium sized Enterprises (SMEs), larger organizations, government agencies and departments and the general public. It also works collaboratively with law enforcement agencies, defence agencies, the UK's intelligence and security agencies and international partners. The NCSC has supported the development of this Code of Practice.

Further details on the operations of the NCSC can be found at https://www.ncsc.gov.uk.

The NCSC has produced guidance on cyber security, which is available at:
https://www.ncsc.gov.uk/section/advice-guidance/all-topics.

Key messages

This Code of Practice is written for safety and cyber security practitioners and their managers to support their understanding of the issues involved in ensuring that the safety responsibilities of their organization are addressed in the presence of a threat of cyber attack. Those engineers have a duty to inform and influence higher levels of management, up to board level of the organization, to bring about the necessary changes. The board and higher levels of management have an overarching responsibility to take ownership of cyber security and safety risks and ensure their staff are suitably qualified and supported.

If safety-related operational technology is not secure, you can't be confident it's safe: absolute safety and security cannot be achieved; the assurance of safety-related systems involving digital technology relies on effective cyber security to reduce the risk of harm to an acceptable level.

The implementation of effective cyber security will in general require modification of the safety-related systems. A close interaction between respective engineers is therefore vital. However, teams responsible for safety and for cyber security are often in different parts of an organization. In many organizations, the governance of the combined risk only comes together at a point of such seniority that the technical competence and capacity for detail may be inadequate to ensure the teams work together effectively. Consequently, the combined risk to the enterprise is not always fully comprehended. **Any divergence or conflict between safety and security goals requires the business to make a conscious risk-based decision on how to proceed.**

Furthermore, the complexity of systems that use digital technology can invalidate the more traditional approaches based on component faults used in safety analysis; and the needs of safety may not be provided by current approaches to cyber security. **The current versions of many standards do not adequately address this relationship** or enable coherent thought about the risks.

Safety and cyber security are mostly complementary risk-based approaches and this Code sets out some shared principles, based on a systems engineering approach, which organizations should adopt and implement within their own context. Engineers will need to work across disciplines, using judgement to manage the risks arising, alongside appropriate traditional disciplines and published procedures.

Implementing this Code of Practice will likely require modification to current procedures. Safety and security need to be considered throughout the systems' lifecycles. However, safety and cyber security standards contain many different expressions of lifecycle phases, which can fail to illustrate how safety and security should work together as reflected in this Code. For example, the adequacy of proportionate security measures needs to be reconsidered as security assumptions change.

This Code of Practice aims to capture the best available practice at time of writing in an evolving topic. It is not intended to provide a means of compliance with any cyber security or safety-related regulations.

Safety and cyber security standards have not yet developed sufficient maturity to provide guidance in detail to organizations on the manner in which these two disciplines should interact within an organization. **This Code of Practice will be revised as good practice for this topic develops.**

Document structure

Figure 0.1 A map of the document and its content

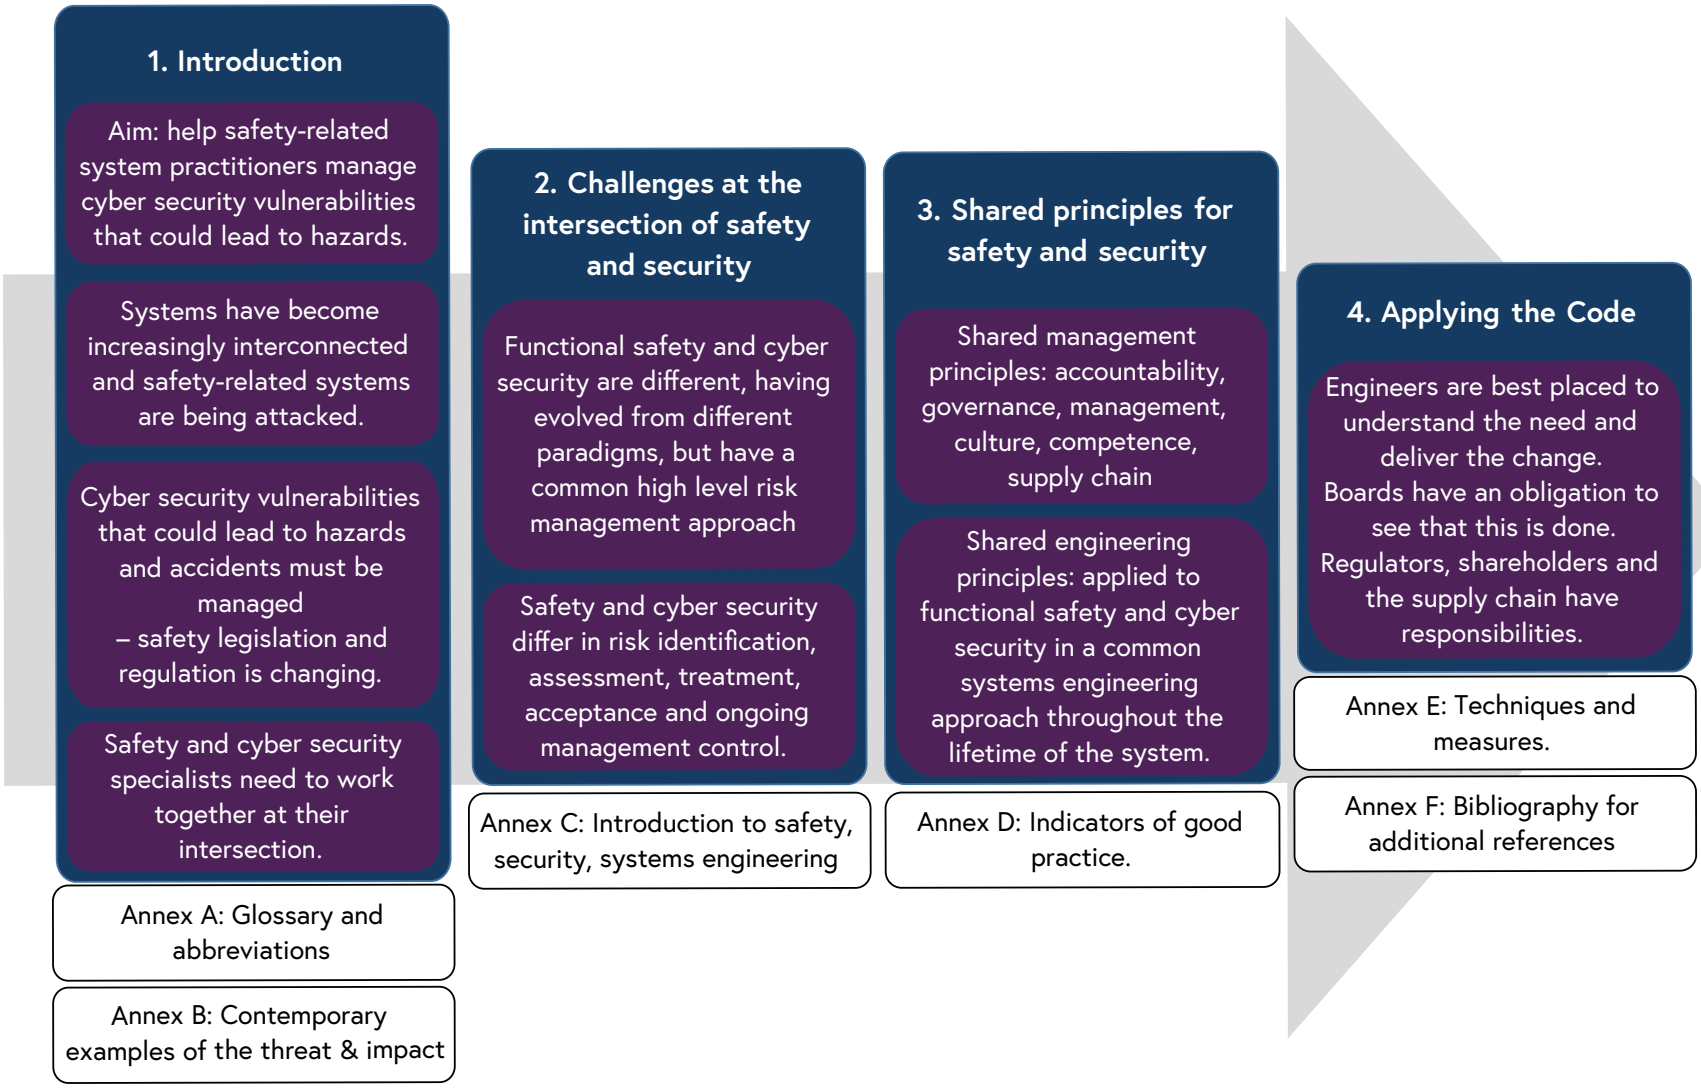

Document structure

Key messages and Section 1 provide awareness of the emerging issues and the need for action in non-technical language that is suitable for consideration by an organization's senior executives.

Section 2 describes the intersection of safety and cyber security and why addressing it can be difficult. This is important contextually, particularly for engineers who predominantly have experience either in safety or in cyber security.

Section 3 is the heart of the Code of Practice and sets out a framework of management and technical principles. Each principle has a brief explanation of why it is important, good practice that addresses the principle and guidance for determining an appropriate approach within an organization.

- Section 3.2 sets out management principles – accountability, governance, management, culture, competence and supply chain – that may lead to specific actions for risk owners and dutyholders who have accountability and the ability to influence policy, strategy, resources and process/procedures.
- Section 3.3 sets out the technical principles – systems engineering, proportionality, risk management and through-life aspects – that may lead to specific actions for system architects, designers, operators and maintainers. The technical principles may also be used by assessors of new and existing systems to guide judgement of whether the issues have been adequately addressed. Note that the supply chain is also addressed by the technical principles.

Section 4 identifies the stakeholders who will have a role in delivering the changes called for by this Code of Practice, with policy and governance led from the top of the organization.

Annex A through to Annex F provide further information.

Section 1

Introduction

1.1 Aim and objectives

The aim of this Code is to help organizations accountable for safety-related systems manage cyber security vulnerabilities that could lead to hazards. It does this by setting out principles that when applied will improve the interaction between the disciplines of system safety[1] and cyber security[2], which have historically been addressed as distinct activities. The objectives for organizations that follow this Code are:

(a) to manage the risk to safety from any insecurity of digital technology[3];
(b) to be able to provide assurance that this risk is acceptable; and
(c) where there is a genuine conflict between the controls used to achieve safety and security objectives, to ensure that these are properly resolved by the appropriate authority.

It should be noted that the focus of this Code is on the safety and security of digital technology, but it is recognized that addressing safety and cyber security risk is not just a technological issue: it involves people, process, physical and technological aspects. It should also be noted that whilst digital technology is central to the focus, other technologies can be used in pursuit of a cyber attack. For example, the study of analogue emissions (electromagnetic, audio, etc.) may give away information being digitally processed and thus analogue interfaces could be used to provide an attack surface.

Many of the issues and proposed principles and practices may have applicability to the wider safety and security disciplines. This Code focuses on the functional safety impact where cyber (in)security contributes, but much of the text has broader utility in safety and security synergies, such as process/governance, etc. It does not advocate the complete integration of the disciplines, but does advocate addressing them in a coherent and balanced manner as part of an integrated systems engineering discipline.

This Code addresses the enterprise level, but the scope may also have applicability to the management of individual products and services.

1.2 Why this subject matters

This subject is important because society relies upon complex digital systems and there is a need for more resilient solutions. It is recognized by some that "if it's not secure, it's not safe" [Ref 1]; this is complicated by the drive for ever greater digital connectivity or 'digitization', the blurring of the boundary between safety and control systems and the increasing complexity of digital systems and their supply chains. Safe and resilient operation of these systems requires that security and safety risks be managed to minimize potential harm in the broadest terms (to people, the environment and organizational assets). Past assumptions that safety-related systems can be operated in a 'trusted environment' are recognized as no longer valid and an adverse environment should always be assumed.

---

1 References to 'safety' throughout this Code may indicate a wider applicability, but are focused on the safety of systems implemented using digital technology, where those systems may be a source of safety risks, or may be used to control safety risks arising from other technology systems. An overview of what is meant by system safety is contained in Annex C, Section C.2.

2 References to 'security' throughout this Code without the cyber prefix may indicate a wider applicability, but the focus of this Code remains on cyber security aspects. An overview of what is meant by cyber security is contained in Annex C, Section C.1.

3 The term 'digital technology' is used to identify the nature of software-based systems that may be most readily attacked. Networking such systems may make them easier to attack and exploit, but attention should also be given to other attack vectors that are not necessarily networked, e.g. Commercial Off-The-Shelf (COTS) devices, programmable/configurable hardware, users that may directly attack systems using replaceable media (maliciously or as an unwitting agent of an attacker), etc.

Section 1 – Introduction

There is awareness of the need to improve cyber security affecting organizations' business information technology[4] assets. However, there has been limited awareness of the impact of poor cyber security on operational technology (OT), particularly as a result of the incorporation of high-performance networked digital technology, thus leaving digital systems exposed to threats of malicious[5] action from cyber attacks.

That awareness is now improving. Recent cyber attacks[6] have demonstrated the potential consequences for essential public services and the impact on individual safety that could be realised by organizations with digital systems that fail to take appropriate measures. Safety regulators[7] are promoting the management of the consequent cyber security risks where these have a potential to result in harm.

Standards such as IEC 61508 [Ref 2] and RTCA DO-356A/EUROCAE ED-203A [Ref 3] recognize that a loss of security may lead to a safety consequence. However, in general, safety and cyber security standards do not guide organizations in how these two disciplines should interact within an organization. Frequently, the responsibilities for addressing safety and cyber security lie in different teams, reporting in to different parts of the organization. This often exposes gaps and may lead to conflicting team objectives and incompatible solutions. These conflicts must be identified and the business must make a conscious decision on how to address them. It is also important to ensure that these considerations are addressed through life, from concept to disposal. Proactive monitoring and response preparedness are essential tools for both safety and cyber security.

An example of apparent conflicts can be seen in civil aviation[8]. Whilst this is not a cyber security example, it serves the point:

A long-standing safety control in civil aviation is to have two pilots on the flight deck of large civil airliners. This mitigates against incapacitation of one of the pilots, for example, due to illness or food poisoning. It also aids workload sharing during abnormal conditions and provides the ability to cross-check correct operation of safety procedures during normal and abnormal conditions. Various measures are in place to reduce the likelihood of common-cause incapacitation, etc.

Following the terrorist hijackings in the United States of America on 11th September 2001, security controls were introduced that led to cockpit doors designed to resist forcible intrusion by unauthorized persons. Access controls could be overridden from inside the cockpit to reduce the possibility of a crew member being coerced into revealing access codes, etc.

In 2015 a suicidal co-pilot locked the captain of an Airbus A320-211 out of the cockpit and deliberately flew the aircraft into the ground. Following this, additional safety controls were put in place to provide better attention to pilots' mental and medical health, and to require two crew members to be on the flight deck at all times. It has been noted that the 2015 accident is not the first time that a pilot has caused the deliberate crash of an aircraft, and that some of the previous incidents have occurred despite there being a second person in the cockpit.

Other options have been considered, such as to allow remote piloting of an aircraft in an emergency situation, or to remotely control the door from the ground, although concerns about the technology itself and the threat from cyber attack have prevented the suggestions being taken seriously.

It is clear in this example that there is no conflict in the objectives of safety and security to protect the passengers and others that may be affected by the control of the aircraft, but there is conflict in the controls employed and a lack of a perfect solution. There is therefore a need to address the risks holistically and seek an optimal solution, which may need to adapt as the threats change over time.

4 Often also referred to as ICT – Information and Communications Technology.
5 Even indiscriminate malware such as a propagating worm was originally written with malicious intent, though not all victims were intended targets.
6 Annex B provides a brief summary of the threat to operational technology at the time of publication.
7 Including the UK Health and Safety Executive (HSE) and the UK Office for Nuclear Regulation (ONR).
8 https://www.skybrary.aero/index.php/Flight_Deck_Security

1.3 The Code of Practice is based on shared principles

This Code provides organizations with a set of shared principles of good management and engineering to achieve complementary and mutually supportive safety and cyber security. These principles build on good practice in safety and cyber security, exploiting synergies and addressing perceived conflict, such that the safety of systems that are affected by cyber security can be assured. It focuses on the interaction between safety and cyber security. Whilst the principles are derived from good practice across systems engineering, and the safety and cyber security disciplines, this Code focuses on how the two disciplines interact where safety and cyber security need to work together, in particular, where inadequate cyber security poses a risk to safety, as illustrated in Figure 1.1.

Figure 1.1 Scope of safety/security intersection

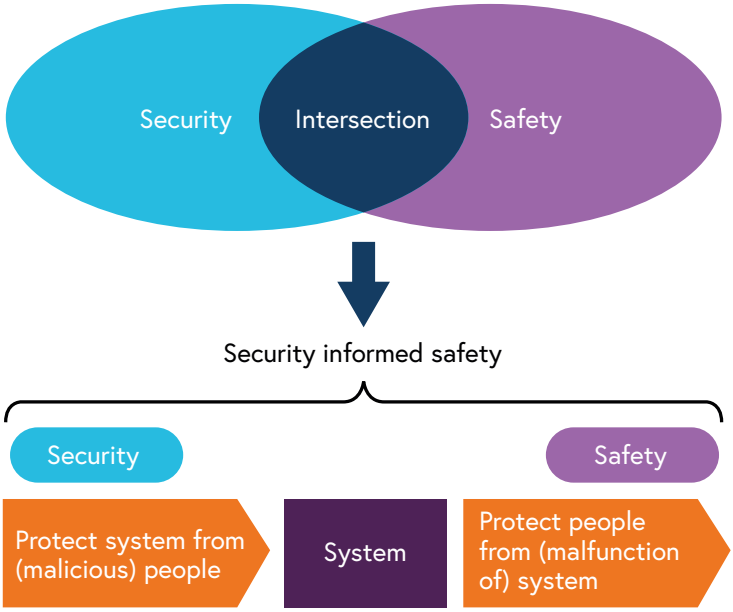

This Code provides recommended practice derived from the shared experience of members of the authoring and technical committees, supported by the study of available literature cited in Annex F. These recommendations are considered to be informative supplements to help address the normative requirements of standards such as IEC 61508 [Ref 2] and ISO 27000 [Ref 12].

This Code can be applied across all domains and internationally, but since terminology used across domains and nations can vary, a glossary is included in Annex A. It seeks to avoid inventing new terms; rather it applies the most natural fit to the topic based on recognized ISO/IEC definitions where practical.

This Code also offers a set of techniques that may be helpful in adhering to these principles and indicates where work remains to be done to develop further techniques. It is intended to be domain-agnostic; that is, it is generally applicable to all domains where an organization needs to address safety and cyber security objectives.

1.4 Engineering judgement will be needed

This Code uses a goal-based approach and certain elements may be unsuitable or insufficient for use in prescriptive regimes. It is not intended to be prescriptive and no single technique/approach will be appropriate in every situation. Addressing the guidance in this Code in a 'tick-box' fashion can prevent organizations questioning the applicability to their situation and lead to over-confidence in how well issues have been managed. When selecting a technique/approach to be applied to a particular cyber security/safety context, it is not sufficient to justify that selection by pointing to its presence in this guidance. Rather, each organization needs to understand its use of digital technology in systems important to safety and consequently judge whether a particular technique is appropriate, and why.

1.5 Using a systems-based approach over the lifetime of the systems

This Code asserts that protection against security compromise should start with the concept of the overall system, through the choice of technologies, architectures and components, to ensure no soft targets are exposed to potential attackers. It also asserts that a prevention strategy alone will be inadequate, and that techniques such as resilience, monitoring and response through life will be necessary.

Whilst today attacks may be perceived to be difficult and uncommon, it can be expected that over the lifetime of a system, a malicious community may commoditize exploits to enable those with much lower skills to attack insecure systems, perhaps because newly discovered vulnerabilities in digital technology haven't been patched quickly enough, or because the systems are inherently insecure through poor design. Consequently, the combination of safety and cyber security should be addressed together in the system design and reassessed during the system lifetime as vulnerabilities are discovered or the nature of the threat environment is judged to have changed.

Consideration is given to in-service systems, legacy systems and off-the shelf elements in addition to new systems at the point of development. Consideration is also given to the entire system lifecycle from concept to disposal. For legacy systems, the same high-level principles apply as for new systems, but the detailed approach may vary. Large systems often comprise and incorporate smaller systems, which may each be at different lifecycle phases.

Particular attention is paid to the interactions through the supply chain, including between operators, integrators, developers, maintainers, suppliers, and throughout the system's lifecycle, including initial acquisition, maintenance and modification, whether by the original equipment manufacturer or a third party.

Adhering to the principles set out in this Code should lead to proportionate, complementary and mutually supportive safety and cyber security measures that are designed, reviewed and revised as necessary throughout the lifecycle of a system.

Section 2

Challenges at the intersection of safety and security

At face value, there is much in common in the consideration of safety and cyber security[9]. Both are dealing with the minimization of risk of an undesired outcome. It is increasingly recognized that security must be taken into consideration when determining the safety of a system, and that potential undermining of safety should be considered when addressing cyber security. There are, however, multiple challenges when the activities of these two disciplines intersect, as described in this Section. There is an urgent need for sharing knowledge, understanding, experience, techniques and approaches between the safety and cyber security communities.

2.1 Reasonably practicable risk reduction is not easily defined for cyber security

There is a recognition that absolute risk elimination is an unachievable goal and so risk criteria are used to determine when an acceptable level of risk has been achieved. It is common practice to require safety risks to be reduced to as low as is reasonably practicable[10,11]. However, there is currently no guidance published about how to demonstrate that these criteria have been satisfied in the context of security protection against cyber attack. Cyber security risk acceptance criteria are often driven by the definition of a risk appetite, set by the risk owner and without any obligation to reduce risk once that appetite has been achieved. Therefore, it is challenging to demonstrate that the measures for cyber security of a safety-related system are sufficient, particularly because of the unknowns about – and rapid changes to – threats and vulnerabilities. Applying current cyber security guidance should be regarded as a minimum, with the safety risk owner making a judgement about whether this is sufficient[12]. This challenge for cyber security is similar to that facing safety when considering systematic issues arising from software and complex hardware.

2.2 Perceived conflict of objectives and expectations

Safety and cyber security communities are often operated in separate parts of the business and it is often perceived that the objectives of the two communities can conflict. For example, the cyber security objectives may require communications to be encrypted to ensure that the end-to-end cyber security of the command and control path is maintained so as to prevent loss of control, or to ensure confidentiality is maintained. However, when viewed from the perspective of the safety objective such encryption may be considered to be introducing an undesirable overhead in terms of communication latency, or to be introducing another function that could fail, leading to unavailability of the safety function.

Each community is familiar with the concept of making design trade-offs (for example, on cost or on system usability) to arrive at practicable system designs and a security/safety trade-off should be seen in the same way as a security/security trade-off or safety/safety trade-off. In some circumstances, safety and security cannot be traded against each other and wider design system changes will be necessary, possibly increasing cost or reducing functionality. The mutual interest in achieving safety and security objectives should make it clear that these issues are not absolutes and that one without the other fails to meet the business needs.

---

9 Annex C provides introductions to cyber security, safety and systems engineering.

10 In the UK, legislation requires that health and safety risks are reduced "so far as is reasonably practicable". The UK HSE uses "As Low As Reasonably Practicable" (ALARP) as the means to determine whether the legal obligation has been satisfied. 'Reasonably practicable' involves weighing a risk against the trouble, time and money needed to control it. Thus, ALARP describes the level to which the UK HSE expect to see workplace risks controlled. [https://www.hse.gov.uk/managing/theory/alarpglance.htm]

11 The 'reasonably practicable' aspect of risk reduction is often seen as a UK peculiarity; however, it is evident in Australian and New Zealand legislation, and the concept is included in the NASA System Safety Handbook.

12 In a regulated environment/industry, it may be appropriate to discuss judgements relating to major hazards/accidents with the regulator.

2.3 Lack of common language

Whilst safety and cyber security have some common objectives, there are frequently issues caused by the lack of common language and frame of reference. Even where the same word is used by both disciplines, it may be defined or interpreted differently, leading to confusion.

2.4 Different engineering perspectives

Security and safety engineers will typically approach a system assessment from different perspectives. Security engineers are more predisposed to think in terms of what a malicious actor might do, including what may appear to a safety engineer as irrational but is calculated to facilitate an attack. Safety engineers are more predisposed to focus on the functions of the system and how those could be incorrectly designed or fail, rather than how they could be manipulated if accessed (that is, accidental or purposeful change to the design/implementation). It has been common historically for safety engineers to be instructed explicitly to exclude consideration of malicious action.

2.5 Quantitative vs qualitative assessment of likelihood

Safety engineers are familiar with reasoning about the likelihood of an unsafe outcome for a defined system and operational context. Some quantitative methods and pseudo-quantitative approaches have been adapted to allow consideration of systematic contributing factors for complex/programmable systems[13].

These approaches are already being challenged as complex software-based systems become more pervasive. Assessing the quantitative reliability of software-based systems is challenging, with the introduction of new technologies, virtualization and software-defined architectures, and behaviour becomes less predictable with the introduction of new technologies (for example multi-core processors, machine learning, etc.), creating greater uncertainty. This is further challenged by the need to express safety risk in the presence of insufficient cyber security controls and an adverse operational environment. Whilst safety has historically used engineering factors to provide safety margins, these are problematic in the face of systematic contributing factors and deliberate acts by an attacker.

Whilst it is always desirable to make risk-based decisions based on quantitative assessment, the nature of complex digital technology and cyber security threats make quantitative assessment of safety and cyber security risks problematic.

2.6 Porous system boundary

The boundary of the system-of-interest used for safety analysis should include those involved in the operation and maintenance of the system[14], but historically has excluded malicious action[15]. This boundary should now reflect the use of digital technology and include the supply chain, for example, where operational maintenance may be remote and cloud-based services may be used. The nature of cyber security risks means that the boundary of the system subject to analysis may have to be expanded to include the malicious actors on the system, whether targeted or incidental, 'insiders' or

13 In IEC 61508 and IEC 61511 the quantitative measure of risk reduction provided by the safety function and the risk of the system failing is derived purely from hardware reliability. The complex hardware and software contribution is premised on the systematic controls applied during the safety lifecycle. If the systematic controls are applied in compliance with the objectives and requirements of the standard, the complex hardware and software is deemed to support the Safety Integrity Level (SIL) of the safety function.

14 The inclusion of these roles means that the system-of-interest may be called a socio-technical system.

15 Malicious action can also be taken to consider malicious inaction by an insider.

external undisclosed third-parties. It also has to consider modifications to the system, such that the system boundary may change, for example, by the introduction of new connections or equipment.

When attempting to analyse the safety effects and likelihood of a system failure in light of a cyber attack, the system boundary becomes potentially unbounded, for example, where the attack involves changing the configuration of the system and potentially supplementing the system with unauthorized devices that permit wider access. Events that may have been considered incredible in a pure safety analysis could become credible due to a cyber attack. The resolution to this dilemma may include restricting the use of some tools and technologies.

2.7 Dynamic nature of cyber security; static nature of safety analysis

The likelihood of a successful cyber attack is significantly more dynamic[16] than that considered in a traditional safety analysis[17]. For cyber security events, a sequence of events required to successfully perform an attack may initially be considered to require significant technical skills and specialist knowledge that make such an attack extremely unlikely.

State-sponsored cyber attacks may be considered possible, as they would be assumed to possess the required resources, subject to having the motivation to apply those resources to the attack. Over time, intellectual property may become more available, vulnerabilities may become known and alternative attack vectors beyond those considered by the analysts may be conceived by the attackers. This can change the practicability of an attack from being state-sponsored to 'script kiddie'[18] and dynamically alter the perceived likelihood of a successful attack. These transformations can occur in very short periods of time. Consequently, it can be hard to justify static assumptions about the likelihood of cyber attacks.

Typically, a safety analysis will consider the likelihood of a given cause to be static, except as determined by usage (scenarios, operational context, demand rates, wear-out) and design change. Monitoring of in-service faults is used to aid detection of differences between predicted and actual failure rates. These differences may be due to differences in demand rate, stress levels, environment etc. experienced by the components. However, whilst historical data is useful for validating hardware/mechanical failure rates, it is unhelpful as an accurate predictor of future cyber attacks or complex failure modes from digital technology, as these are affected by influences outside of random chance.

2.8 Tensions over maintenance of cyber security

Safety-related systems, particularly those with the potential to affect the highest severity outcomes, are expensive to develop and assure, and consequently are often designed to be in service for many decades. Vulnerabilities may come to light long after the original design team have moved onto other products or organizations. Rapid change to address a vulnerability has the potential to introduce other unintended safety consequences. This leads to a culture conflict, where cyber security demands a 'fix immediately' approach, whilst safety demands a 'fix after careful consideration' approach.

16 The concept of 'dynamic likelihood' reflects the change in assessed likelihood over time, given additional information.

17 Safety analysis may need to adapt to address likelihood in a more dynamic manner, due to emerging technologies such as Artificial Intelligence and Machine Learning.

18 'script kiddie' refers to non-serious hackers who shortcut most hacking methods in order quickly to gain their hacking skills. They may use hacking programs written by other hackers, because they often lack the skills to write their own. Script kiddies are considered to be inexperienced and immature, but can inflict as much computer damage as professional hackers [https://www.techopedia.com/definition/4090/script-kiddie].

2.9 Different risk paradigms

The dynamic nature of cyber security and the challenges in assessing potentially unbounded systems make it difficult to establish whether safety risks initiated by all potential cyber attacks have been reduced to an acceptable level. Whilst both safety and cyber security domains consider risk, it is clear that they are considered in different paradigms and so it is not easy to reflect these in a simple two-dimensional matrix.

2.10 Different risk control philosophy

Safety hazards are often addressed using controls that are assessed as independent of each other, whereas a malicious attacker may deliberately exploit more than one system cyber security vulnerability in a damaging sequence of events, which is designed to react to and nullify the defensive measures initiated by the system. The use of redundancy and diversity techniques in safety-related systems may improve cyber security, but may simply present a larger attack surface to a malicious actor.

A common misconception is that cyber security can simply be added as a 'protective shell' around an existing system, without altering that system. This model of protection may be inspired by the model of physical protection provided by locked cabinets, rooms and buildings. Whilst this may be useful as one element of providing cyber security, it is undermined by the use of digital technology, with its intrinsic complexity and porous boundary, exploited by ongoing supply chain relationships, software updates, etc. In many standards, cyber security measures are assigned to categories of protection, detection and response, reflecting the contemporary assumption that any sufficiently motivated and capable adversary will defeat the protective measures and the system will enter a different state while detection and response measures are enacted. Safety and cyber security measures should be designed to work together to detect, respond to and recover from hazardous conditions and/or cyber attacks, despite not being able to describe in advance the nature of all possible cyber attacks.

Section 3

Shared principles for safety and security

3.1 Introduction

The following Sections introduce management and engineering principles that are relevant to addressing the safety/security intersection. The principles are also tabulated in Table 3.1 below, with indications of good and poor practice that may help organizations assess their own performance set out in Annex D.

Table 3.1 Shared principles for safety and security

Principles are expressed as predicate statements that can be evaluated to true/false and are expressed in the current tense. It is intended that these are assessed throughout the lifecycle. Actions to ensure they are addressed may result in practices that produce plans, expressed in the future tense, whilst actions that assure whether they were adequately addressed at a given lifecycle milestone may result in practices that produce reports, expressed in the past tense.

Typically, the Sections follow the structure of an introductory paragraph setting out the background and relevance; a statement of the principle; recommended good practice; and further guidance addressing the principle. The guidance focuses on how consideration of the safety/security intersection may expand on or adapt the good practice already applied in the safety and/or cyber security domains. Throughout all of these, it is recognized that addressing safety and cyber security risk is not just a technological issue, but involves people, process, physical and technological aspects.

3.2 Management principles

3.2.1 Accountability

Risks can arise to organizations from cyber threats that include the health and safety of workers and of members of the public, environment, business operations, security of supply, personnel data, reputation, finance and commercial risks. Many of these will be subject to legal requirements with which the organization needs to comply. Organizations where the risks arise will be responsible for managing those risks and will therefore need to understand what those risks are and set their own risk appetite for each of the types of risks to which they are exposed, subject to any overriding legal requirements, for example, that specify limits to acceptable risk.

The jurisdiction's legislation sets obligations for organizations to address the safety of their operations, their employees and those affected by their activities. It also sets obligations for security and, whilst many of these historically have been focused on confidentiality of personal data and national security issues, legislation [Ref 11] has been introduced to address the continuity of operations that are critical for the national infrastructure[19]. These later aspects are not directly related to safety, but they form part of the context for consideration of security-related legislation that could apply to safety-related systems.

It needs to be recognized that, in addition to the impact on people and the environment, significant safety or security incidents can have a harmful effect on an organization's reputation with customers and shareholders. In addition, potential action may be taken by regulatory bodies to gain assurance on behalf of the public and other stakeholders that an organization complies with legislative duties.

Principle 1:	**Accountability for safety and security of an organization's operations is held at board level[20].**
Practice 1.1:	The board should put in place traceable delegation of responsibility and authority for addressing safety, security and their interaction.

The board is the ultimate owner of the risk. Accountability cannot be delegated, though it is a necessity of business that the means of addressing the accountabilities of the organization are shared through delegation to individuals/groups with clear objectives and that they are empowered with the necessary authority and resources to address their scope of responsibility in a manner that is proportionate to the potential harm.

Whilst an organization may rely on its partners throughout its supply chain to provide solutions to the technical challenges of its operations, and this may include derived requirements to help address safety and security, the organization retains accountability for safety/security performance of the organization. Further information regarding the supply chain is addressed in Section 3.2.5.

Practice 1.2:	The board should require regular and proactive reporting of issues that affect safety/security performance.

Reports should be unbiased and include both positive and negative aspects as appropriate, using suitable metrics and identifying trends and key issues. Reports should enable the board to understand the risks and, where appropriate, to take action to ensure there is clear guidance to those with responsibility on tolerability of risks, to ensure that appropriate resources are being applied and that there are no gaps or ambiguity in responsibility for addressing the risks. Reporting criteria should be established to ensure that reporting does not overwhelm the organization's ability to process the reports and make appropriate decisions: that is, to ensure that the right information gets to the right people at the right time.

19 Operations that are critical for the national infrastructure are known as 'essential services' under the terms of the Network and Information Systems (NIS) Directive [Ref 11].

20 Specific accountability may apply in some sectors, for example, the 'Licensee' in the Nuclear sector, or an 'Accountable Manager' in Aerospace.

3.2.2 Governance

An integrated approach to systems engineering requires a holistic approach to governance of the organization. This requires that the organization has identified the relevant legal and regulatory frameworks within which it operates, and that the duties that those frameworks place on the organization are addressed in the business policies, processes and procedures.

Principle 2:	**The organization's governance of safety, security and their interaction is defined.**
Practice 2.1:	The board should set clear policies for safety and security.

Policies should be periodically reviewed to ensure they remain fit-for-purpose relative to threat landscape, risks, organizational structure and maturity, business strategy, etc. Appropriate behaviour throughout an organization is enhanced if the board takes an active interest in monitoring the performance against these policies. This can take the form of ensuring that suitable resources are made available, that proactive reporting of leading indicators of safety/security performance is a regular part of briefings to the board, that good performance is rewarded, and that constructive corrective action is taken where deficiencies are noted. The board is key to establishing a culture where shortfalls can be reported and addressed prior to the shortfalls escalating to a serious incident (see also Section 3.2.3 on culture). Learning from incidents and shortfalls, and sharing this learning across the organization, can help to facilitate such a culture.

Practice 2.2:	The policies should encourage safety and security to be addressed co-operatively as part of an integrated systems engineering approach[21].
Practice 2.3:	The organization should establish governance mechanisms to identify synergies and resolve conflicts between the objectives specific to safety and security.

Where responsibility for different risk aspects is through distinct reporting lines, for example, operational safety via the operations director and cyber security risks via the information technology director, the policies need to make clear how the interdependencies are addressed. It may be necessary to change organizational structures or responsibility boundaries to achieve this. The information technology director may not appreciate the responsibility for operational technology cyber security, meaning that product security aspects are often overlooked. The differences between cyber security of information technology and that of operational technology mean that it may be better to separate the operational technology cyber security responsibilities and allocate them to the operations director.

The governance mechanisms should identify reporting and escalation criteria and mechanisms to address cases where any conflict cannot be resolved by those with delegated responsibility.

Principle 3:	**Demonstrably effective management systems are in place.**
Practice 3.1:	The organization should operate management systems that require the identification of relevant legislation and regulation.

In addition to legislation and regulation, the management systems should identify how the risk criteria are set. This may be constrained by the legislation and regulation, but will also be influenced by the organization's risk appetite and policy.

Good practice for safety includes operation of a Safety Management System (SMS) that defines the approach taken by the organization for recognizing and addressing its responsibilities. Similarly, good practice for cyber security would include operation of a Cyber Security Management System (CSMS)[22]. These management systems have considerable overlap in the approaches required and both in turn rely on management of factors such as quality, documentation/information and assets governed by related management systems.

21 Systems engineering is not limited to the design/development phase, but should be applied to all phases, including operation in service.

22 Akin to an Information Security Management System (ISMS).

It should be noted that existing management systems may lack full coverage of the issues raised in this Code. For example, the standard Information Security Management System (ISMS) is not designed to handle the complexity of many organizations involved in the design and operation of safety-related systems. Specifically, the boundary of the systems often excludes operational technology and the ISO 27000 [Ref 12] approach does not readily accommodate security of suppliers, a particular issue in industries with complex supply chains supporting the operational technology.

Practice 3.2:	The management systems should be designed to ensure that they identify inter-dependencies and interactions to ensure compatibility.
Practice 3.3:	The Safety Management System and Security Management System should be parts of a comprehensive and coherent high-level management system.
Practice 3.4:	The management systems should be maintained to ensure they manage safety[23] and security risks using current relevant good practice.
Practice 3.5:	The management systems should include measures to detect shortfalls against safety and security objectives and also identify and address the cause(s) of such shortfalls.

The governance arrangements will need to address assurance that the management systems are effective. For example, the arrangements will need to ensure that safety and security performance meet applicable legal requirements and the board's risk tolerability policy.

Principle 4:	**The level of independence in assurance is proportionate to the potential harm.**
Practice 4.1:	The organization should establish criteria for the use of independent assurance against safety and security objectives, including the scope of independent assurance activities and level of independence.

A common approach to achieving confidence in assurance justifications is to employ independent resources to do one of the following:

(a) conduct assurance activities such as verification and validation; or
(b) perform the risk assessment; or
(c) audit the activities of the organization in addressing their safety/security objectives.

This may be achieved through the appointment of an agent by a customer, through the procurement of consultancy services or within the organization itself. Consideration could be given to establishing an assurance department within an organization that is are independent of the delivery streams and reports directly to the board, similar to Independent Nuclear Assurance and Airworthiness departments, but with scope to address safety and cyber security. Applying the same scope and level of rigour to all classes of system would unduly burden low risk operations. Since there is significant uncertainty in factors of likelihood, the potential level of harm should be taken as the leading indicator to determine a proportionate approach. Other factors that may be taken into account include complexity and novelty. Some standards, for example, IEC 61508 [Ref 2], include requirements and guidance on the level of independence required.

The assurance should address safety and security risks in a holistic manner, paying particular attention to the interdependencies between safety and security. Independent assurance between independent non-communicating safety and security experts is inadequate. A team of independent assessors with suitable level of expertise in safety and security could be used to conduct the independent assurance activities.

23 In the UK, safety legislation requires risk of harm to be reduced so far as reasonably practicable. This includes where risk of harm may arise from security threats.

3.2.3 Culture

Even in the most mature and effective management systems, mistakes will be made that lead to shortfalls in performance against their objectives, and incidents may occur. An organization that can acknowledge these and put in place measures to detect them and address the root cause(s) will reduce the risk of a serious outcome that could have been reasonably foreseen and prevented. Causes may include tensions between the objectives of the management systems that have not been fully recognized/managed. Whilst processes and procedures help, the culture of the organization is more important in promoting an atmosphere where shortfalls can be freely identified and reported, such that corrective action can be taken before the shortfall escalates to a full incident.

Principle 5:	**The organization promotes an open/learning culture whilst maintaining appropriate confidentiality.**
Practice 5.1:	The organization should ensure that reporting of shortfalls against safety and/or security objectives in a trustworthy and responsible manner is encouraged.
Practice 5.2:	The organization should ensure that the emphasis of investigation into shortfalls is seen as learning to avoid future shortfalls, rather than to allocate blame.

A learning culture needs to avoid perception of unreasonable punishment for reporting of a shortcoming. Management systems and the processes/procedures that implement them should actively seek the counterexamples that show there is shortfall in performance. Counterexamples may be identified through incidents and accidents, or through anomalies discovered after they have escaped the process checks that should have detected and resolved them. Having identified these, actions to identify causal aspects and opportunities for intervention should be pursued with a view to future avoidance rather than blame. Study of shortfalls across the organization may help to identify trends that are not evident in isolation, and present opportunities to learn from a different part of the organization, even in the absence of an issue locally. Learning should also include an exercise to extrapolate shortcomings as creatively as possible to imagine unexpected implications.

Practice 5.3:	The organization should ensure a proactive approach to learning, through continuous improvement, training and sharing with the wider community.

These learning philosophies can be extended to all aspects of the business (including the supply chain), through setting incentives to identify and address shortfalls before a critical incident occurs. An open channel with the supply chain to understand the implications of newly discovered vulnerabilities or attack vectors is particularly important. This can be even more powerful if the learning is able to draw on experience of a wider user base[24].

3.2.4 Competence

The rapidly evolving nature of technology and its impact on safety and cyber security risks is such that legislation, regulation and standards often do not attempt to specify prescriptive methods to address the risks. Legislation, regulation and standards are generally goal-based and rely heavily on the competence of those involved in development, operation, assurance and governance to interpret their objectives and provide solutions that satisfy those objectives.

Using recognized competency frameworks can assist with managing appropriate competency within an organization by providing both structure and competency levels. Frameworks may not contain the specific criteria required by an organization, but they can be supplemented. For example, the IET has published a *Code of Practice for Competence for Safety-Related Systems Practitioners*; this does not explicitly cover the competencies required to address security of a safety-related system and there is no direct equivalence for cyber security. Further information is in Annex E.2.

24 For example, the NCSC runs the Cyber Security Information Sharing Partnership (CiSP), which is a joint industry government confidential forum for sharing intelligence about cyber threats.

Principle 6:	**Organizations are demonstrably competent to undertake activities that are critical to achieving security and safety objectives.**
Practice 6.1:	The organization should identify the key competencies required to achieve their safety and security objectives.
Practice 6.2:	The organization should identify how the competency requirements are allocated across their organizational structure to groups and individuals with accountability, responsibility and/or authority for the setting or achievement of safety or security objectives.

The competencies required to address the intersection of safety and security adequately necessitates a unique set of skills and experiences. It would be uncommon to find this set within an individual. The competencies can be addressed by bringing together individuals with relevant competencies to operate as a cohesive team on the appropriate tasks such that shortfalls in one individual's competencies are balanced by the strengths of another. For example, this could include inviting a cyber security engineer to a Hazard and Operability Study (HAZOPS) exercise, or a safety engineer into a cyber security review, or could require a larger multi-disciplined team.

The team does not have to be established in an organizational structure, but it does need to be formed for the relevant tasks with a clear purpose and with a clear line of authority and reporting of team outcomes. Care needs to be taken to ensure that the separate management reporting lines of individuals within the team do not conflict or distract from the shared objectives of the task.

Competency therefore has to be considered as a combination of organizational, team and individual characteristics and needs active management to establish suitable culture, practice and governance.

Digital technologies, the systems that use them and the cyber threats they face are continuously evolving. It is therefore necessary to ensure that competencies are up to date for the tasks being conducted. This may be achieved through refreshed training and exercises to ensure practical experience.

Practice 6.3:	The organization should record how it has ensured that the competency requirements are satisfied, and how these are maintained over time.

Careful consideration is required of how key competency and knowledge records are managed and attributed to individuals. Special measures will be required to treat records of dismissals and disciplinary actions relating to personnel in positions of authority in the interest of protecting the organization and the rights of the individuals.

3.2.5 Supply chain

To achieve its objectives, any sizeable organization operation is likely to rely on a supply chain for system elements, services and resources. The supply chain can provide access to expertise to supplement the organization, but can also be a source of cyber security vulnerabilities and attack. An overarching organizational strategy of how to address safety risks and protect against cyber attack is required. This will result in allocation of requirements and responsibilities to the supply chain, and also to the technical, procedural and commercial interfaces with the supply chain. An organization may also choose to outsource cyber security and/or safety assessments and advice. It is important that it has enough core competence in house that it can comprehend the assessments and advice provided and retain responsibility for its decision-making.

Principle 7:	**The organization manages its supply chain to support the assurance of safety and security in accordance with its overarching safety/security strategy.**

The procurement/acquisition process ensures that a product or service provides the required cyber security and safety functions and services while meeting all concerns and constraints expressed in the requirements. Accountability for the performance of safety and security cannot be delegated to the supply chain. This does not prevent the organization holding the supplier to account contractually for the satisfaction of the derived requirements.

Practice 7.1:	The organization should assess the nature of the relationship it needs with its suppliers in order to meet enduring obligations to provide cyber security services (e.g. patching, incident response support, etc.) for the lifetime of their products and services.

There are significant challenges in securing the required relationship for extended durations. The organization may not have sufficient commercial or market presence to be able to impose the desired controls on a supplier, and the supplier's business may go through several transformations through acquisition or even collapse over the lifetime. Only by assessing the potential needs for a long-term relationship and identifying the challenges can the organization devise a strategy for how to resolve these.

Practice 7.2:	The organization should identify the requirements and responsibilities allocated to their supply chain to support achievement of the safety and security objectives.
Practice 7.3:	The organization should identify the controls and reports it will use to manage the supply chain to ensure appropriate oversight of the factors that impact safety or security, including those that arise from emergent functionality/behaviour of the supplied product.
Practice 7.4:	The organization should ensure it has sufficient in-house competence and capacity in any outsourced safety or cyber security assessment and advice services that it can retain control of its risk decision-making.

The controls used by the organization to manage the supply chain are part of their management systems. In order to assess the risks that exist at the operations level, it is necessary to understand how the immediate suppliers address their allocated responsibilities, including how they manage these through their own supply chain. It is also necessary to understand what protection and mitigation is in place at the organizational level should the interfaces be compromised.

This should also place requirements on the cyber security of the organization, and its development environment. The adoption of approaches such as Cyber Essentials [Ref 13] and ISO 28000 [Ref 30] in supply chain contracting can help, but should be seen as the minimum and cannot be expected to achieve the same level of risk reduction as a well-considered and targeted set of measures to identify key controls and responsibilities.

3.3 Technical principles

These principles are grouped as 'technical' in that they are addressing the risks related to the technical/engineered systems[25]. It is recognized that addressing safety and cyber security risk is not just a technological issue, but involves people, process, physical and technological aspects. The responsibility for addressing these principles will depend on the nature of the organization and it remains the responsibility of the management of the organization to identify policies, allocate responsibilities and provide resources to facilitate effective consideration of these principles.

The principles need to be applied in all scenarios, whether 'greenfield' development or in-service operations, and whether using bespoke design, re-use of legacy designs or Commercial Off-The-Shelf (COTS) equipment.

The topics are set out in a logical sequence starting with system scope definition and moving through the steps of risk management to risk acceptance and management through life. It is emphasized that these should not be seen as a single pass linear process. The principles are intended to be applied continuously throughout the life of a system.

25 It is noted that the UK NCSC published a number of secure design principles that may prove useful in addressing some of these technical principles. See Section F.2.3 (Web resources).

3.3.1 Systems engineering

Safety and security are emergent properties of systems. Like quality, it is good practice that these areas should be addressed as part of the system's whole-life engineering activities rather than addressed separately and 'bolted-on'. An introduction to systems engineering and how it relates to safety and security is contained in Annex C.

It may be useful to consider systems by their role in the organization. IEC 62264-1 [Ref 14] uses a functional hierarchy model that can be abstracted and built on for more general use. Figure 3.1 illustrates that some functions may be more associated with traditional information technology systems, whilst those closer to the low-level product of the organization may be distinguished as operational technology. Safety-related systems typically exist at Levels 0/1/2. It should be recognized that this is a functional model and will not represent the physical or data model unless the systems are deliberately designed to achieve this.

Figure 3.1 Information/operation technology functional hierarchy

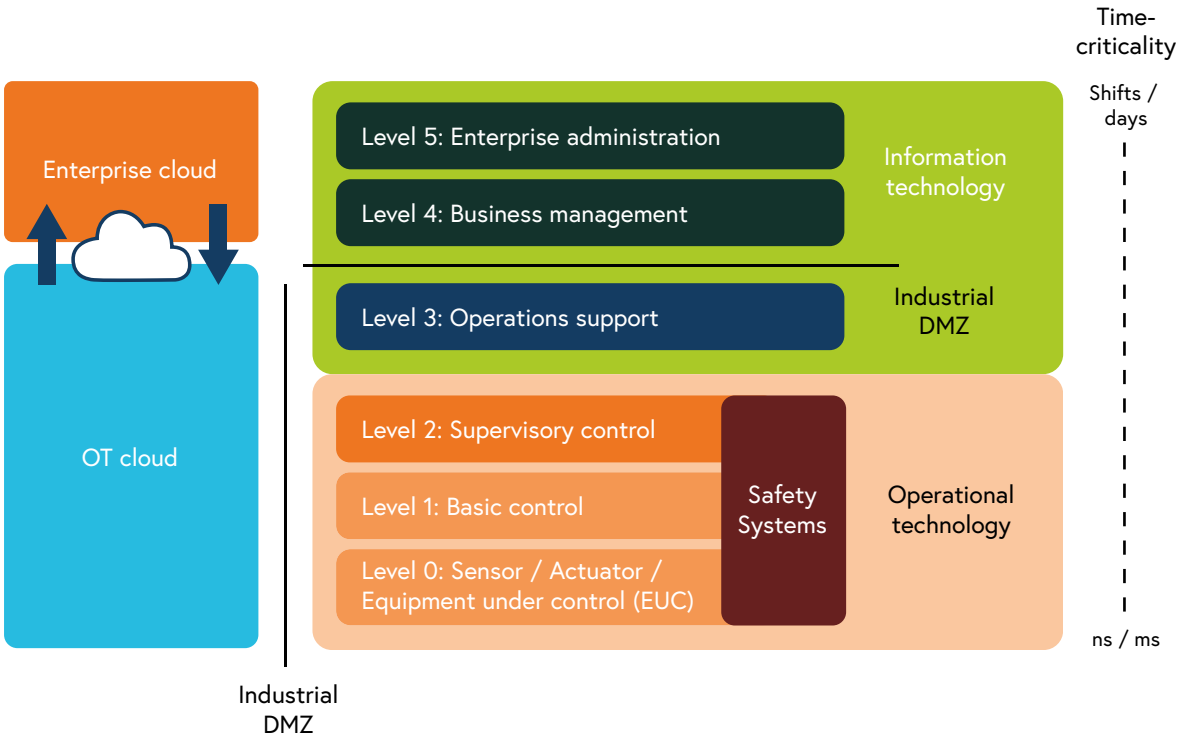

In this model, Level 5 may be focused on external stakeholders such as shareholders, potential customers and the general public, involving communications such as email and internet access, as well as communications from finance and human resources (HR). Such systems will typically receive much attention from business information technology security but are unlikely to be directly relevant to a safety concern. In the context of the ISO 15288 [Ref 15] system context model, these are likely to be 'systems in the environment' rather than systems-of-interest.

The lower level functions are those most likely to be related to safety hazards as a cause or control. They could include safety-related control systems such as a Flight Control System in an aircraft, or protection systems that work independently from the primary control system, such as a Safety Instrumented System (SIS) in an industrial control process. These are likely to form a system-of-interest. An organization will need to consider the overall operational system as its system-of-interest, whilst the perspective of an organization in the supply chain will move to adopt one or more of the sub-systems as their system-of-interest. A system-of-interest to an organization in the supply chain will have its own enabling systems and interacting systems in the environment.

Principle 8:	**The scope of the system-of-interest, including its boundary and interfaces, is defined.**
Practice 8.1:	The organization should identify, document and communicate the scope of the system-of-interest within the bounds of its safety and security objectives.

A decision should be made about what is included in the nominal system-of-interest. The term 'nominal' is introduced to encourage consideration of how the system scope may be affected by malicious activities.

Practice 8.2:	The organization should identify interacting systems in the operational environment and the enabling systems that could impact safety or security objectives.

It is important to identify interacting systems in the operational environment and the enabling systems such that their interfaces and potential involvement in an attack can be addressed. It is useful to consider the engineered interfaces as well as the inherent interfaces with the environment that could be exploited by an attacker.

Principle 9:	**Safety and security are addressed as co-ordinated views of the integrated systems engineering process.**
Practice 9.1:	The organization should define engineering and business processes that enable separate security and safety disciplines to co-ordinate their activities against the safety and security objectives.

Safety and security are recognized as separate disciplines, with their own skill sets and considerations, but the overlap in system considerations and interdependencies mean that they should not be addressed independently. Safety and security practitioners need to understand clearly which outcomes are possible, and which are not, from a cyber compromise of the safety system. Conflicts need to be addressed proportionately to the consequence.

Organizations may choose to pursue a tighter integration of the activities that support safety and security processes. A number of previous studies of the safety and security of software-based systems have suggested there are significant underlying similarities and that there are efficiencies to be achieved by addressing the two in a fully integrated method, for example, SafSec[Ref 16]. Whilst these potential advantages in integration of methods are recognized, care has to be taken to recognize the differences in skills and knowledge required to apply the methods to achieve the different safety and security goals. More recent work has provided a basis for co-assurance against safety and security standards. [Ref 31] proposes a Safety-Security Assurance Framework (SSAF) to provide a systematic approach to reasoning about safety and security.

The National Institute of Standards and Technology (NIST) has published two volumes on systems security engineering. Security topics are addressed in the context of the system lifecycle processes. The volumes are designed as complementary guidance:

1. *Systems Security Engineering – Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems* [Ref 23] provides an engineering perspective and describes the actions necessary to develop more defensible and survivable systems, in light of the growing adverse consequences of cyber attacks, disruptions and hazards, where the need for trustworthy secure systems is paramount.
2. *Developing Cyber Resilient Systems: A Systems Security Engineering Approach* [Ref 24] can be used as a handbook for achieving required cyber resiliency outcomes from a systems engineering perspective on system lifecycle processes, utilizing the experience and expertise of an engineering organization to determine what is correct for its purpose.

3.3.2 Proportionality

It is generally recognized that safety and security risks cannot be completely eliminated in practice, and that there will always be some level of residual risk in order to achieve some desired benefit. It is therefore necessary to address the issue of how much investment in terms of time and resources should be applied to addressing the risks. Individually, safety and security domains have established approaches to help resolve this issue, but there are difficulties when addressing it as a combined risk.

Principle 10:	**The resources expended in safety and security risk management, and the required integrity and resilience characteristics, are proportionate to the potential harm.**
Practice 10.1:	The organization's safety and security management systems should define frameworks and criteria that guide engineering activities to achieve a proportionate approach.

Traditional frameworks and criteria used by each discipline in isolation may need to be refined to address the interaction of cyber security with safety objectives. In the safety domain, proportionality considerations can take into account the severity of an outcome and the likelihood that such an outcome will occur. There are significant complications with estimating the likelihood of a successful cyber attack and therefore, when considering the proportionality in the context of a cyber attack on a safety-related system, an indirect proxy for likelihood, such as difficulty in execution of a successful attack, may be considered more appropriate. It should also be recognized that the level of difficulty can change rapidly and therefore caries a high degree of uncertainty across the lifetime of a typical safety-related system. Where there is significant uncertainty in likelihood assessments, it may be appropriate to consider just the severity of harm that could occur and, where this is significant, to apply the precautionary principle[26]. A proportionate investment in the security controls required to mitigate the risk of a cyber security breach would be relative to its potential effect on the uncertainty of the likelihood or severity of a safety system risk.

Some safety legislation allows risk reduction to cease at the point where further effort would be grossly disproportionate to the reduction in risk. For safety-related systems, this may also be applied to security controls that protect safety functions. Given the difficulties in determining what measures can be considered reasonably practicable (see Section 2.1), the approach to be taken should be discussed with the appropriate regulator.

Other factors that should be taken into account in considering proportionality are complexity and novelty. Both of these are matters of perspective and need to be evaluated with respect to the organization's experience as well as wider industry references. For example, increased levels of uncertainty, complexity or novelty may require greater levels of independence from assurance organizations as required by IEC 61508 [Ref 2].

26 It should be noted that the precautionary principle applies in the case of scientific uncertainty, does not imply zero risk, and is not an excuse to ignore taking action/decisions based on the uncertainty. For a discussion of the precautionary principle, see [Ref 28]. In the European Union, the Precautionary Principle has a legal definition and application, see [Ref 29].

3.3.3 Risk management

Risk management comprises a number of activities[27] that are applied iteratively through life, as depicted in Figure 3.2:

Figure 3.2 Risk management process cycle
(Source: derived from ISO 31000:2018 *Risk management – Guidelines*)

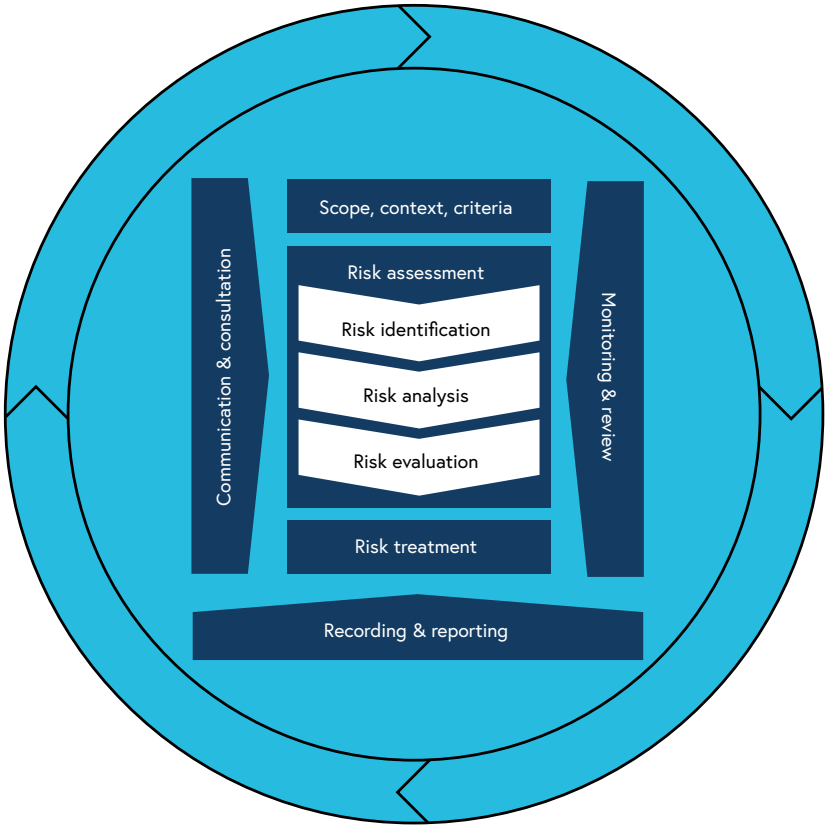

Risk management of both safety and cyber security addresses people, process, physical and technical aspects of the system. It is applied iteratively, throughout the system's lifecycle. All elements are important aspects of the risk management process and should be addressed in the organization's risk management processes. This Code highlights a number of considerations that are relevant to the intersection of cyber security and safety.

3.3.4 Risk assessment

Risk assessments are always context dependent and therefore rely on an understanding of the operational use and environment. The criteria for reassessing risk should be defined in the relevant management policies and plans and should include relevant changes in context of system use or design, following relevant incidents and at a maximum period in the absence of other criteria being met. Relevant incidents are not restricted to those that are directly associated with safety or security, but should also include any that may challenge or support assumptions or context used in those risk assessments.

Risk assessment involving systems comprising complex electronics and software is difficult. Traditional techniques originally devised for relatively low complexity systems often did not adequately address the risk arising from systematic errors or human factors, or restricted their attention to elements that could be readily assessed quantitatively. Approaches to address these aspects continue to evolve.

27 These form the core of the risk management process as defined in ISO 31000 [Ref 17]. This model is chosen as a unifying underpinning concept for both safety and cyber security risk management cycles, each of which have their own nuances, but can be related back to these core concepts.

Principle 11:	Safety and security assessments are used to inform each other and provide a coherent solution.
Practice 11.1:	Techniques should be selected that are appropriate to the lifecycle point and the objective of the assessment.

Different techniques may be applicable at different lifecycle phases and when assessing at different levels of abstraction. Risk assessments can be component- or system-driven and this is true both in a cyber security context or a safety context. This can help in identifying techniques that can be applied at each lifecycle stage and level of abstraction. No single technique is suitable in all cases and a clear understanding of the objectives of assessment and the strengths/weaknesses of each technique are important for selecting an appropriate complement of techniques.

Practice 11.2:	Interaction points between security and safety assessments should be identified and communicated.

Many of the techniques, whether they come from a safety or security approach, can be adapted to enable a complementary assessment of safety in a security threat environment. Further details are discussed in Annex E. It can also be useful to apply standard safety and security assessment techniques without adapting them, but ensuring that interaction points are clearly identified and communicated. This allows a safety team to identify critical areas of concern, and a security team to assess how a security attack may compromise those areas.

The objective of risk assessment is to determine the level of risk that exists for a given operation. It may be used to prioritize resources to manage the risk (see Section 3.3.6) or to determine that the residual risk is at a level that can be accepted (see Section 3.3.7). The detail of a risk assessment will aid in the determination of risk management strategies. Risk assessment may be used as a part of design trade-off analysis or as part of change management prior to the introduction of a design or operational procedure change.

Practice 11.3:	The results of assessment, whether quantitative or qualitative, should be expressed together with a measure of confidence.

Safety engineers are familiar with applying a range of qualitative and quantitative safety assessment techniques in order to determine appropriate derived safety requirements, and to assess the level of residual risk left to be managed. It is acknowledged that quantitative assessment is impractical early in a system development, and that quantitative evaluation of risk is difficult even for a mature system design, particularly for complex software-based systems, and for human interaction. Similarly, it is difficult to ascribe quantitative measures to aspects of a successful security attack. Given these challenges, an integrated qualitative[28] approach that considers both safety and cyber security risks is likely to be required.

The approach taken for assessing the safety risk in the presence of a security risk, and whether to use qualitative, quantitative or semi-/pseudo-quantitative[29] assessment, will depend on the available information and the purpose of the assessment. It is important to recognize the confidence in any quantification. The suitability of the approach used and the confidence in its conclusion should be justified.

28 A longer-term objective may be to have a quantitative approach to risk assessment, and this is an area of active research by the academic community. The current state of the art of quantitative assessment of security lacks maturity, and there are challenges for its application to the safety of complex digital technology. To be usable, such an approach would have to be able to make meaningful predictions of the future likelihood of complex socio-technical events. It is therefore not advocated by this Code.

29 Semi- and pseudo-quantitative encompass the use of terms that represent a range of values (such as 'frequent', 'unlikely', etc. to represent a probability/frequency range) or terms that are an abstract representation (such as high/medium/low), which have no definitive quantification but are used to represent a likelihood/probability trend.

Where there is an absence of trustworthy quantitative data, it may be appropriate to base risk decisions such as prioritization or proportionality on the potential severity of undesirable outcome, rather than attempting to factor in likelihood. This is likely to apply in the early stages of deciding how much effort to place on security risk avoidance for a safety critical function.

At some point it will be necessary to determine whether a system design has achieved a level of risk sufficient to enter or continue operation. At that point, at least an implicit measure of likelihood is being applied, for example, an assertion that security measures are sufficiently likely to protect against attack that the safety risk assessment evaluation can stand (see Risk acceptance – Section 3.3.7).

3.3.5 Risk identification

The objective of risk identification is to identify, as early as possible, the key undesirable outcomes[30] associated with the system that may arise during the life of the system. Risk identification is an important part of project planning, requirements definition and business case investment justification. It also helps target where more thorough analysis should be undertaken. A successful risk identification and assessment will help to gauge the effort that is likely to be required to provide an appropriate design solution and to support an effective assurance argument, proportionate to risks. It provides input to:

(a) identifying any critical areas of risk inherent in the organization's requirements;
(b) providing the supporting evidence for any investment case;
(c) scoping the subsequent activities required;
(d) selecting or eliminating investment options for subsequent assessment;
(e) setting the initial safety and security requirements;
(f) providing the starting point for subsequent analysis; and
(g) initiating safety/security risk tracking and management.

Principle 12:	**The risks associated with the system-of-interest are identified by considerations including safety and security.**

A broader scope for consideration of risks may be conducted, for example, to include wider stakeholder concerns such as loss or unforeseen degradation of mission, reduction of capability/capacity, or other business concerns.

Practice 12.1:	Risk identification should be performed using team-based techniques employing a broad range of competencies and viewpoints.

Consideration should be given to the use of a team of people that includes safety specialists, security/cyber specialists, operators and maintainers from within the organization and potentially wider industry, academia or government organizations.

Practice 12.2:	Techniques that provide a systematic basis for risk identification should be employed.

A number of techniques are identified in Annex E.

Risk is defined in many standards as being a function of severity and likelihood. During risk identification, only the undesirable outcomes and the impact/severity of their consequence should be assessed. The likelihood cannot be ascertained with any confidence at this point, particularly in light of any security threat, and may mislead the prioritization and proportionality of effort in the subsequent activities.

30 Undesirable outcomes relate to the loss of something of importance to stakeholders. Depending on the domain, these may be referred to as accidents, incidents, mishaps or adverse events. Typical losses of relevance to safety and security include loss of life or injury to people, loss of availability (of a safety function or mitigation) and loss of confidentiality.

Risk identification is initially performed early on in a system lifecycle to provide the earliest possible opportunity for treatment of the exposed risks. However, it shouldn't be seen as a single opportunity to identify risks. New risks may come to light as the system, its concept of use and the understanding of its interaction with other systems and its environment matures. It is also necessary to consider interaction of risks, for example, cascade risks, where the realisation of one risk may lead to exposure to other risks – this is at the heart of the consideration of the cyber security effect on safety risks.

Practice 12.3:	All identified undesirable outcomes, together with assumptions or identified interactions, should be recorded before being rationalized.
Practice 12.4:	Risks should be rationalized to remove duplication and produce a coherent set of risks to be addressed.

A common issue with the listing of risks and hazards is that they narrow the focus onto a particular perception of the risk that addresses only a particular cause. This can narrow the mindset of the analyst, thus obscuring other sources or alternative solutions.

Practice 12.5:	The process of rationalizing risks should identify system hazards at a consistent level of abstraction.

Conventions such as those recommended in the System-Theoretic Process Analysis (STPA)[31] handbook [Ref 18] can help to manage these issues, including stating risks without reference to individual system elements or causes and stating hazards in terms of a system state or conditions that are within the system's control without reference to the external environment.

Practice 12.6:	Care should be taken not to rationalize risks on the basis of lack of a credible cause, or the extreme unlikelihood of the cause occurring.

Security attacks are potential sources of causal events that may be considered incredible through component failure or operator error alone. Furthermore, a targeted attack may deliberately affect multiple elements of a system to create co-ordinated causal events that may not be considered credible through component failure and/or operator error. They may intentionally mask the discovery of a failure/hazardous state by inhibiting or manipulating diagnostic reporting/annunciators.

3.3.6 Risk treatment

For complex safety-related systems, safety standards and risk approaches have been successfully applied for a number of decades, but are based on the need to manage faults and their effects. Faults can occur due to errors in specification or implementation, or through defects occurring in service through component failure, maintenance or operator errors. Faults can be latent (that is, present, but not revealed until some trigger condition or event) or exposed (that is, resulting in some erroneous performance or behaviour). When considering security, the source of faults has to be extended to include accidental or targeted malicious attacks that modify the implementation. Vulnerabilities can be seen as a form of fault. Potential faults are managed using techniques that fall into four main categories:

1. fault avoidance – correct by design, the use of mathematical proof, etc.;
2. fault removal – through test, inspection, etc.;
3. fault tolerance – resilience, avoiding single points of failure, etc.; and
4. fault warnings – allow time for recovery/emergency procedures to reduce the risk of harm.

31 This Code recommends a mix of techniques, providing both top-down and bottom-up views. STPA is an example of one of these techniques, but its handbook provides some useful guidance that can be applied generically. Its application in the context of this Code is described in Annex E, Section E.3.

Risk control measures are applied in order of precedence according to a scheme such as that illustrated in Table 3.2:

Table 3.2 Risk control precedence example

Measure	Example
Elimination	Redesign the system(s) so that the risk is eliminated. This is particularly required for hazards that can result in societal harm and significant loss of life.
Substitution	Replace process with a less hazardous and less vulnerable one.
Engineering controls	Design in additional controls to protect data and information. Reduce the energy that could result in harm. Minimize the exposure (number of people or amount of time at risk). Increase number of barriers that restrict propagation to a dangerous state. Give priority to measures which protect the system collectively over individual measures, e.g. anomaly detection, intrusion tolerant architectures. Introduce measures to reduce systematic defects.
Administrative controls	These are all about identifying and implementing socio-technical procedures to maintain the integrity of the system to address known faults and are the least effective controls.

The objective of applying these measures is to eliminate the risk where practicable, or to lessen the severity and/or likelihood of an undesired outcome. There are clear parallels between safety and security risk treatment. Whilst the first priority is to eliminate vulnerabilities, measures that make it more difficult to exploit potential latent vulnerabilities will reduce the likelihood of an attack successfully leading to the undesired outcome.

When selecting these measures, it is necessary to consider whether the solution introduces any new risks, hazards or causes of hazards. It may be necessary to perform a trade-off evaluation to determine an optimal balance where it is impractical to reduce one risk without impacting another. For example, in safety-critical applications, system performance (response time, latency in communications, etc.) is often an important non-functional property that may be affected by additional security controls such as encryption. Different aspects of security may also be in conflict.

As it is impractical to eliminate all risks, there is a significant reliance on measures that enable unknown flaws, including vulnerabilities, to be tolerated without impacting on safety or business objectives. Such tolerance should be at least long enough for the attacks against the vulnerabilities to be detected, responded to such that any risk is avoided in the short term and recovered from such that normal operations and risk management processes can continue.

Principle 13:	**System architectures are resilient to faults and attack.**
Practice 13.1:	System architectures should be engineered to employ multiple layers of protection that make attacks more difficult, more likely to be detected and responded to and less likely to lead to harm.

The typical safety management approach to engineering resilience is to use redundancy and diversity in the implementation of the system architecture to provide defence in depth. When considering the security view, it has to be acknowledged that a targeted attack is likely to attempt to defeat such measures. There is also a view that the introduction of additional assets to achieve redundancy/diversity simply increases the surface for security attacks. Whilst an attacker may still have to compromise multiple systems to cause a safety event, it may be easier for them to achieve some other malicious objective such as denial of service, or a confidentiality breach. Such successful attacks could influence the effectiveness of the safety measures by retrieving information that makes future co-ordinated attacks more effective, or by instilling a lack of confidence in the system by operators, making them more likely to take decisions that undermine the safety objective (for example, overriding safety protections to avoid perceived 'nuisance' faults). A motivated attacker could retrieve information over an extended time period

of several months or years as part of a reconnaissance activity so that they can simultaneously defeat multiple layers of protection. It may therefore be necessary to adopt more rigorous security measures for redundant/diverse system elements than may have been considered based on their individual purpose.

Practical examples of measures that have been used successfully in the nuclear sector include:

(a) the introduction of an independent cyber resilient layer within the architectural model that does not contain computer-based equipment but relies on non-complex logic (for example, magnetic logic), enhancing both safety and cyber security.

(b) the introduction of measures (for example, data diodes) to prevent data communication from lower integrity systems. This approach can improve robustness of independence arguments and eliminate data pollution.

Practice 13.2:	Maintenance and operations activities should have an enduring obligation to ensure that the engineered safety and security protections are not compromised or circumvented.
Practice 13.3:	Systems and components should be designed such that security controls can be maintained and updated, in operation, in a safe manner.

The engineered resilience needs to be assured through life and throughout maintenance and modification. This requires a record of the design basis[32] and a discipline to ensure changes are assessed before implementation. System audits may be required to ensure that operator/maintainer action has not introduced unauthorized changes.

Practice 13.4:	Components should be hardened against targeted attacks, according to their role in fulfilling the safety and security objectives.

Systems-based assessments will help to identify strategies to avoid the key hazards/losses, whilst component-based assessments can help to assess the implementation for robustness of the defences against targeted attack. Ideally, components will be selected that have been hardened against attack at source (Secure by Design). The selected components should come with manuals that describe the assumptions made regarding use of the product within systems and specify how to operate and maintain the component securely and safely. This may include installation/integration requirements and constraints to maintain the safety and security properties of the component in the system context.

An important part of resilience is the ability to detect an anomalous state/event and the preparedness to respond to and recover from that state/event. This establishes expectations for system monitoring, planning, and learning/sharing, which should be addressed to detect a fault and control its effects before it propagates to system hazard. Safety management should ensure that incident/accident planning identifies abnormal/emergency operating procedures. To address the cyber security threat, these procedures need to be adapted to identify the precursors to an incident that may exist with a security attack, recognizing that a targeted attack may deliberately obscure warnings and indicators that would otherwise indicate a hazardous fault in the system. Such an attack may be part of a co-ordinated targeted attack, or incidental to an attack that has no specific target within the affected system. Security intrusion monitoring may aid detraction of precursors.

Practice 13.5:	The possibility that mitigation and recovery systems may be affected by a security attack should also be considered.

Response to a hazardous event may include operating in a constrained or degraded mode until safe recovery can be achieved. It is not uncommon for safety-related systems to be designed to 'fail-safe' on detection of an anomalous condition. This can be exploited by a security attacker to affect a denial of service where their objectives are not necessarily to achieve a harmful effect. Resilience considerations when addressing combined safety and security risks may therefore direct a more considered approach.

32 Design basis includes the design configuration baseline together with the requirements baseline, concept of operation/use, rationale for design decisions, etc.

Practice 13.6:	Organizations should clearly identify and communicate the events and activities that lead to a decision to move the safety system into or out of reduced functionality mode (to protect the integrity of the safety system against attack).

Recovery from a security-related incident may include restoring a compromised system to a known good state and cleansing it of operational data. It may also require modifications to the system to prevent a recurrence of the attack. Care has to be taken to ensure that such modifications do not compromise either safety or security objectives. Prior planning that pre-empts an attack will enable operators to take appropriate action in the event of a detected attack, and subsequently to recover to normal operations when the threat has been dealt with. Such transitions between normal and reduced states, and vice versa, should be assessed and trained for to ensure they are safe and do not expose vulnerabilities in the transition. 'Design for change' strategies will help enable modifications to be introduced and assured with minimal consequential impact.

Such prior consideration is especially important for organizations that have to maintain live operation. Business continuity and disaster recovery plans are an important part of ensuring appropriate response to a range of scenarios. Default 'fail-safe' strategies may be the simplest to implement, and appear reasonable when considered purely from a safety viewpoint, but may not be necessary where adequate levels of safety can be maintained using other strategies. They may be undesirable from a more holistic viewpoint, may create other safety risks in the wider environment and may be exploited by attackers to create a disproportionate effect.

3.3.7 Risk acceptance

Risk acceptance is one possible risk treatment outcome. It should be noted that risk acceptance does not mean that no further action is required. In order to 'accept' a risk, as a minimum, it is necessary to maintain the measures that have been specified, and the assumptions made, that make the risk acceptable. It is also necessary to re-evaluate whether the context in which that acceptability has been determined remains valid. This should take place at a pre-determined periodicity, but should also follow any significant change in use or operating environment and should occur in response to incidents.

Principle 14:	**The risk justification demonstrates that the safety and security risks have been reduced to an acceptable level.**
Practice 14.1:	Risk criteria should be established by the organization that bound tolerable levels of risk against safety and security objectives.

Safety standards typically require safety criteria to be established that enable residual risks to be accepted. These will typically include a risk matrix that will identify how severity and likelihood are combined to provide a measure of risk, and this is then mapped to tolerability and authority levels. The criteria will often also require a tolerable level of risk to be achieved and that risks are reduced where reasonably practicable.

With the inclusion of security considerations, it is also necessary to provide criteria for the effectiveness of security measures. Risk matrices applied to security risks can be extremely misleading, particularly given the issues in assessing likelihood, and have led to poor risk management decisions. This may be based on a consideration of how a security risk modifies the likelihood of a safety outcome occurring, or based on a confidence measure of how effective the security measures are at making the security impact on safety protection negligible. There is a lack of generally recognized criteria that can simply be adopted and it is particularly challenging to apply traditional safety criteria. All relevant good practice should be considered in order to demonstrate adequate risk management, and where appropriate, the relevant regulator should be engaged. Once determined, the definition of what is tolerable should be documented, understood, accepted and practised.

Practice 14.2:	The residual risk of harm against safety and security objectives should be justified against the risk criteria.

As the means of achieving an acceptable level of risk is determined on a case by case basis, a simple standards compliance approach is rarely adequate to justify the acceptability of residual levels of risk, or the efficacy of the through-life risk management measures. It is therefore common for standards to require an explicit justification of the approach taken and the results achieved. In some sectors, this is captured in a safety case[33], whilst in others it is provided by a collection of assurance documents. Historically, it has been acceptable for safety assurance documents to make assumptions about security, or to exclude considerations of security from their scope. It is increasingly recognized that security has to be considered and justified in a more integrated manner. An explicit assurance case, covering both safety and security considerations, would extend the reach of the traditional safety case.

Practice 14.3:	A holistic approach should be taken to the justification of system trade-offs, seeking alternatives that provide optimal satisfaction of all required properties, including achievement of safety and security objectives.

Where trade-offs are required of system performance against different safety or security objectives, the level of residual risk against each objective needs to be justified and particular care has to be taken in considering unintended impacts from one to the other. Where practical, the conflict should be resolved such that the sources of risk are removed, for example, by removing network connections added for convenience. Where this is not practical, the trade-off needs to be justified to the satisfaction of the dutyholder and relevant regulator. For example, where the availability of a safety function is reduced by adding a security function, the justification may appeal to the potentially more severe effects of a safety outcome should the absence of the security function lead to failure of the safety function through a security attack.

3.3.8 Through-life management

Considerations of residual risk from a traditional safety approach focus on technical measures taken in development, combined with operational measures applied through life. This is supplemented through life by monitoring and fed back into operation, maintenance and, where appropriate, technical modifications.

The nature of security concerns means that, whilst it is important to address technical measures in development, considerable effort on continuous assurance is required through life. As the security landscape changes, attackers acquire knowledge about products used in the system and adapt their techniques and methods.

Principle 15:	**The safety and security considerations are applied and maintained throughout the life of the system.**
Practice 15.1:	Activities to address safety and security objectives should be planned and addressed at every stage of the entire lifecycle, from concept to disposal.

The management and technical principles set out in Sections 3.2 and 3.3 are intended to be applied throughout the lifecycle, from concept to disposal. The most important effect is that achieved during operations through the in-service phase; however, the considerations against the principles need to be applied through life. Through-life management requires regular consideration of evolving security threats, as well as evolution of the system and its use. Safety and security considerations should form an integral part of the bigger Asset Management System [Ref 32], where one is used.

The disposal phase is often overlooked, but insecure disposal of operational technology can provide a source of information and development/test environments for attackers, helping them develop and perfect the capability to attack similar systems. Disposal should consider skills and knowledge as well as physical assets and documentation.

33 The primary purpose of an operator's safety case is to demonstrate to themselves that the risks associated with their operations meet their safety objectives. The key here is that dutyholder needs to manage the risk; the regulatory function is to assess whether the duty has been discharged. Consequently, in practice, the safety case is needed to run an operation safely rather than to satisfy the regulator.

Practice 15.2:	Activities to address safety and security objectives should be addressed for every aspect from architecture to component level, and from concept to disposal.

Considerations made during architectural design can help to design systems that do not require, or are tolerant to, changes in configuration, for example, by addressing the need for regular patching. It should also be recognized that controls may be different between phases; for example, controls during commissioning may be different to controls during operation.

Risk management should not be seen as a one-off exercise. It should be planned into the lifecycle to support key risk acceptance decisions, consistent with a systems engineering approach. These key risk acceptance decisions include decisions to proceed with a concept into development, accept a product into service and whether risks are becoming too high to tolerate requiring a modification or replacement, or withdrawal of services. Planning of what these decision points are, and when they should occur, helps to ensure efficient and effective management of the risks. Legislative and regulatory obligations to maintain and update risk assessments exist in many sectors. Lifecycle decision points may include investment gates early in a system concept/development phase, transition to operation, maintenance/upgrade points and ultimately decommissioning and disposal.

Practice 15.3:	Planning for risk assessment review should recognize that the evolution of the threat environment is potentially more dynamic than that of the technology and operating environment.

The assurance justification's update cycle needs to be commensurate with both the technology refresh rates and the evolution of the operating environment, including the threat environment. A planned and well-structured assurance approach will aid rapid consideration of the impact on the justification in light of new threat environment information.

Practice 15.4:	Change management should be an active process operating from concept to disposal.

Changes to the configuration of the system are a common source of compromise of safety or cyber security objectives. Planned and authorized changes can have unintended consequences if not carefully considered. For example, even those perceived as very minor, such as the introduction of a link that makes maintenance easier, can defeat the carefully planned separation of critical systems from those that are open to the internet. Unplanned changes can be introduced through poor understanding of the system architecture by an unwitting design, operation or maintenance practitioner, or can be introduced through malicious action.

Changes can occur in the way that operations are conducted through non-technical effects such as changes in processes and procedures, changes of people in practitioner, support or management roles and changes in the culture of the organization. Such changes are often made with the intent of delivering a positive benefit, but can also have unintended consequences for safety or cyber security objectives.

Robust change management processes are required that:

(a) include consideration of people, processes/procedures and culture, as well as technical aspects;
(b) assess authorized changes to ensure they have the required effect, and do not introduce unintended effects that impact negatively on either the safety or cyber security objectives;
(c) minimize the likelihood and impact of unauthorized/unintended changes; and
(d) monitor the system configuration to proactively identify unauthorized/unintended changes.

Section 4

Applying this Code of Practice

4.1 This Code of Practice is written for engineers and engineering management

The topic of this Code is potentially of interest to a wide range of roles within a business organization. However, it is not practical to address the needs of all potential stakeholders in this Code. Therefore, **this Code is primarily focused on engineers and those who manage them (typically an intermediate level of management)** to support their understanding of the issues of ensuring that the safety responsibilities of the organization are addressed in the presence of a threat of cyber attack. This Code is not restricted to those with safety/security in their job titles, rather it is addressed to those engineers and engineering managers with the potential to influence the approach taken by the organization.

In supporting this understanding it is intended that the informed engineer and manager **can influence higher levels of management up to board level of the organization,** such that the importance of the issue can be recognized at levels that can set appropriate policies and governance arrangements, supported by the necessary structure and resources to provide a proportionate and effective response.

It is also intended that the informed engineer and manager can enable changes to procedures and approaches to ensure awareness of practitioners and facilitate practice that manages the cyber security/safety risks in a proportionate manner. **This Code asserts that existing safety and cyber security procedures will need to be reviewed and may need changing.** This activity should cover the breadth of the organization's activities from concept to disposal, and through supporting functions including supply chain management and customer support.

This Code recognizes that safety and cyber security are not just technical issues; they involve people, process, physical and technical aspects. Annex E describes techniques and measures that an organization may wish to copy and interpret to suit their business context, to deliver the principles set out in this Code. Annex F provides a bibliography and additional references.

4.2 Other stakeholders

There are a number of stakeholders with an interest in, and an ability to influence, the successful outcome of addressing the safety-security intersection. A high-level approximation of the organization elements is shown in Figure 4.1. This also illustrates a simplified view of the responsibilities and information flow relating to the management of safety and security risks.

Figure 4.1 Nominal information and management flows
(Source: derived from NIST's *Framework for Improving Critical Infrastructure Cybersecurity* [Ref 10], Figure 2)

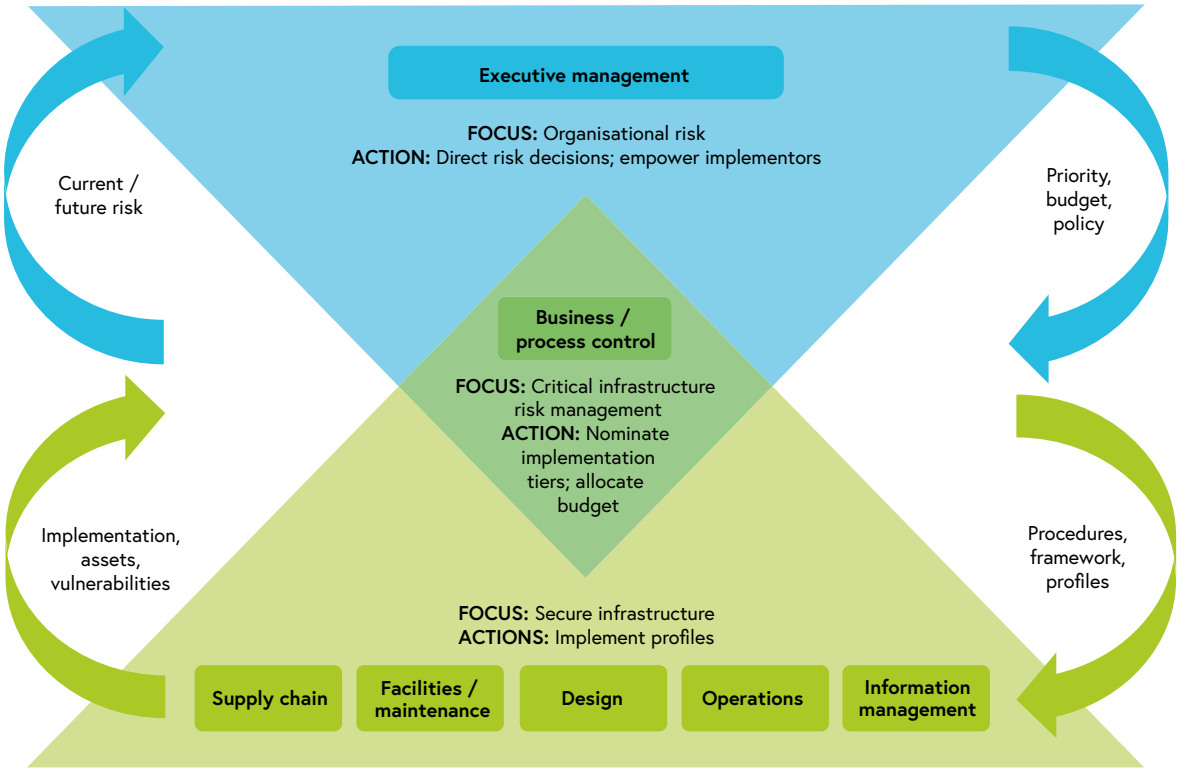

4.2.1 The board

The consequence of a safety incident resulting from a cyber security incident has, by definition, the potential to cause harm to people, property or the environment, but can also have significant business implications, such as the financial costs of recovery, disruption to business and impact on reputation and share value. It is therefore a concern for board level consideration from a commercial, legal and ethical viewpoint.

Setting in place adequate measures to protect the business requires action from the most senior leadership of the organization to address culture, governance and business practices that enable the whole organization, including all elements through to the base of the supply chain, to play their part.

The occurrence of a safety-related incident that has been caused by a security lapse may have an adverse impact on:

(a) regulatory compliance and licence to operate;
(b) business continuity, revenue and cashflow;
(c) reputation and market position;
(d) litigation;
(e) profitability; and
(f) shareholder value.

These can be affected by adverse publicity affecting customer/public confidence and shareholder value; legal action taken on behalf of the affected parties seeking compensation (for actual or claimed harm) resulting in legal fees to defend against the claim; regulator interest imposing penalties and/or increased oversight in ongoing activities; a change in the perception of risk by insurance underwriters leading to higher insurance fees, and so on.

The board has the ability to exercise the greatest influence over the organization through policy, culture and investment. A successful, well-managed intervention in an attempted cyber attack can boost confidence and mitigate or reverse the potential negative effects identified above.

4.2.2 Shareholders

Whilst this group is unlikely to have a detailed interest in the specifics of measures taken to address the security-safety intersection, they will undoubtedly have an interest in occurrences (actual or perceived) that undermine confidence in the ability of the organization to manage the risk where this impacts on share value.

4.2.3 Regulators

Regulatory good practice is to apply a proportionate approach, on the basis of risk. There is a tendency in some countries for regulation of areas relevant to safety and security to adopt outcome/goal-based regulation[34], moving away from prescriptive regulations[35].

This Code emphasizes the use of good practice to deliver appropriate safety and security performance and addressing their interaction which, if implemented correctly, should be acceptable to the regulator.

4.2.4 Other colleagues in management

Those with responsibilities for the management of operations need to understand and interpret policies set at board level for addressing safety and security risks. This will include:

(a) establishing appropriate organizational structures for managing and reporting on the effectiveness of addressing these risks;
(b) the allocation of appropriate resources in terms of budget, time and competent people;
(c) the determination of suitable processes and procedures that address the risks throughout the lifecycle;
(d) the clear identification of accountabilities and responsibilities; and
(e) supporting a culture that encourages reporting of issues and learning from experience.

4.2.5 Supply chain

Suppliers of products and services will have similar concerns to those set out above for their supply element. This is a two-way relationship, where operational incidents can impact on the reputation of a supplier, and where issues in the supply chain can create risks for the organization to manage.

Where the regulatory obligation rests with the operator of the service, it is essential that the operator translates the regulatory requirements to be met by the supply chain vendor into the commercial contract. This should not be taken to excuse the supplier from applying good practice to the design, manufacture and support of its product.

The nature of digital technology may begin to place a new and difficult responsibility on suppliers of products with a software element. They may be obliged to provide patches and ongoing support of their products, as supporting software such as operating systems are updated, for the lifetime of the product, which could exceed 25 years. Furthermore, providers of such products may form an important part of any incident response arrangements. This may not be covered by existing contracts. For critical components, the relationship with the supply chain may no longer be simple and needs to reflect a continuing service relationship.

34 Regulatory expectations for both the HSE and the ONR are provided in guidance used by inspectors, notably [Ref 25] for the HSE and both Appendix 6 – Cyber security of computer based systems important to safety of [Ref 26] and FSyP 7 Cyber Security and Information Assurance of [Ref 27] for the ONR.
35 UK safety regulation has been goal-based since 1974.

Table of Contents for the Annexes

The annexes provide further information to supplement the main body of this Code.

Figure 0.1 Copy of a map of the document and its content (showing the relationship of the Annexes to the content of this Code)

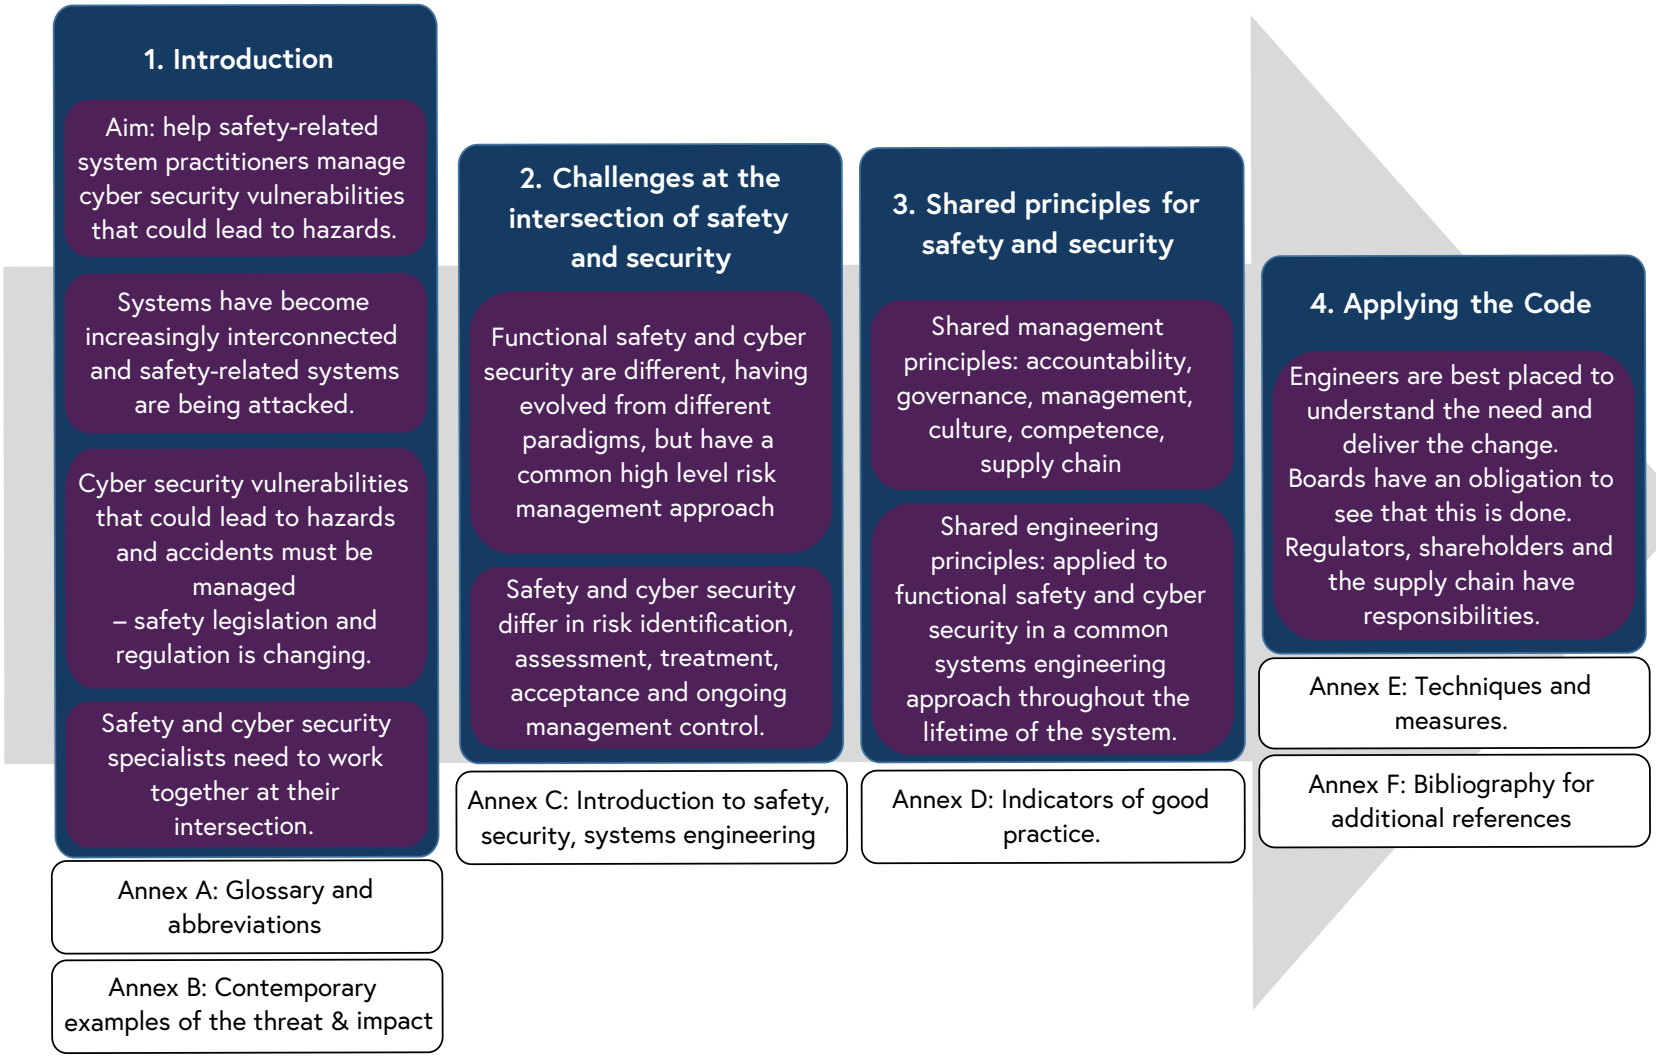

Table of Contents for the Annexes

Annex A

Glossary and abbreviations

A.1 Abbreviations

Abbreviation	Expansion
ALARP	As Low As Reasonably Practicable
ASEMS	Acquisition Safety & Environmental Management System
BCS	British Computer Society
CAF	Cyber Assessment Framework(https://www.ncsc.gov.uk/collection/caf)
CCP	Certified Cyber Professionals
CCSC	Certified Cyber Security Consultancy
CCTV	Closed Circuit TeleVision
CIISec	Chartered Institute of Information Security(https://www.ciisec.org/) [formerly known as 'Institute of Information Security Professionals']
CISA	Cybersecurity and Infrastructure Security Agency
CiSP	Cyber Security Information Sharing Partnership
COMAH	Control of Major Accident Hazards
COTS	Commercial Off-The-Shelf
CSMS	Cyber Security Management System
E/E/PE	Electrical/Electronic/Programmable Electronic
EN	European Norm
EP&R	Emergency Preparedness and Response
EUC	Equipment Under Control
FMEA	Failure Modes and Effects Analysis
FSyP	Fundamental Security Principles
FTA	Fault Tree Analysis
GDPR	General Data Protection Regulation
HAZOPS	Hazard and Operability Study
HMI	Human Machine Interface
HR	Human Resources
HSE	Health and Safety Executive (UK – https://www.hse.gov.uk)
IACS	Industrial Automation and Control Systems
ICA	Independent Cyber Assessment
ICS	Industrial Control System
IEC	International Electrotechnical Commission (https://www.iec.ch)
IET	The Institution of Engineering and Technology (https://www.theiet.org/)
IIoT	Industrial Internet of Things
INCOSE	International Council on Systems Engineering (https://www.incose.org)
ISMS	Information Security Management System
ISO	International Organization for Standardization (https://www.iso.org)
IT	Information Technology
LOPA	Layers of Protection Analysis
MIT	Massachusetts Institute of Technology
MOD	Ministry of Defence (UK – https://www.gov.uk/government/organisations/ministry-of-defence)
NCSC	National Cyber Security Centre (UK – https://www.ncsc.gov.uk)

Annex A – Glossary and abbreviations

Abbreviation	Expansion
NHS	National Health Service (UK – https://www.nhs.uk)
NIS	Network and Information Systems
NIST	National Institute of Standards and Technology (US – https://www.nist.gov)
NUREG	Nuclear Regulation (US – https://www.nrc.gov/)
OED	Oxford English Dictionary (https://www.oed.com/)
OES	Operator of Essential Services
ONR	Office for Nuclear Regulation (UK – http://www.onr.org.uk)
OT	Operational Technology
PLC	Programmable Logic Controller
PPE	Personal Protective Equipment
RITICS	Research Institute in Trustworthy Inter-connected Cyber-physical Systems
RTCA	Radio Technical Commission for Aeronautics (https://www.rtca.org)
SCADA	Supervisory Control And Data Acquisition
SFIA	Skills Framework for the Information Age
SIL	Safety Integrity Level
SLA	Service Level Agreement
SMS	Safety Management System
SSAF	Safety-Security Assurance Framework
SSEP	Safety, Security and Emergency Preparedness
STAMP	System-Theoretic Accident Model and Processes
STPA	System-Theoretic Process Analysis
SWIFT	Structured What-If Technique
SyAP	Security Assessment Principles
TPN	Technical and Professional Network
UCA	Unsafe Control Actions
UK	United Kingdom
US	United States (of America)

A.2 Glossary

The source of these descriptions is indicated in square brackets. Further discussion of the use of the term in the context of this document is in plain text, following the formal description, which is in *italics*.

Term	Meaning
assurance case	*reasoned, auditable artefact created that supports the contention that its top-level claim(s) is satisfied, including systematic argumentation and its underlying evidence and explicit assumptions that support the claim(s) [ISO /IEC /IEEE 15026-1:2019]* The subject of the claim(s) of an assurance case can include safety and/or security.
cyber security	*The process of protecting information by preventing, detecting, and responding to attacks. [NIST Cybersecurity Framework]*
cyber-physical systems	*comprise interacting digital, analog, physical, and human components engineered for function through integrated physics and logic [NIST]*
enterprise	*Entrepreneurial economic activity [OED]* Can be taken as synonymous with a business organization or company. It may be used to imply a top tier of a hierarchy (e.g. enterprise level); however, there is no formal definition that supports this use. Enterprise has become common in the context of information technology (e.g. enterprise class, enterprise solutions). To avoid ambiguity, the term 'organization' is preferred in this Code.
functional safety	In this Code, the term is used in a very broad sense as relating to the safety of, arising from, or controlled by functional elements of a system (typically implemented using digital technology, but not exclusively). It may be thought of as relating to safety issues that arise from what a system does, rather than from what it is. In this sense a system can include people. This usage is broader than often perceived from the IEC 61508 definition: *"part of the overall safety relating to the EUC and the EUC control system that depends on the correct functioning of the E/E/PE safety-related systems and other risk reduction measures" [IEC 61508]* (EUC = Equipment Under Control)
operational technology	*hardware and software involved in monitoring and control of physical devices, processes and events [https://www.gartner.com/en/information-technology/glossary]* 'operational technology' is a term used particularly in Industrial Control and often abbreviated to 'OT'. It is used in its expanded form in this Code to aid familiarity for those outside the Industrial Control industry. With the advent of the Industrial Internet of Things (IIoT) and cloud technology, the boundary between information and operational technology is becoming less distinct.
organization	*An organised group of people with a particular purpose, such as a business or government department. [OED]* Can be taken as synonymous with 'Enterprise'. The term 'organization' is preferred, to avoid ambiguity with information technology connotations of 'enterprise' and to avoid any legal framework associated with 'business' or 'company'.
safety-related system	*Any system whose correct operation is important to ensuring safety.* This includes those systems considered safety-related by IEC 61508, but is more broadly interpreted in this Code.
system element	*member of a set of elements that constitute a system [ISO 15288]* including hardware, software, data, humans, processes (e.g. processes for providing service to users), procedures (e.g. operator instructions), facilities, materials and naturally occurring entities or any combination. A system element is a discrete part of a system that can be implemented to fulfil specified requirements.
system-of-interest	*system whose lifecycle is under consideration [ISO 15288]* The perception and definition of a particular system, its architecture and its system elements depend on a stakeholder's interests and responsibilities.

Annex B

Contemporary examples of threats and potential impact

Examples of recent cyber attacks that had a potential impact on safety are introduced below. Further details are available within the references cited.

- The WannaCry cyber attack that seriously affected the NHS on Friday 12th May 2017 highlighted cyber security risks in the health sector. The widespread disruption to health services was well publicized. The Committee of Public Accounts report HC 787 [Ref 4] assessed the financial implications of the disruption to service, whilst the press[36,37], reported on the potential safety implications of delayed treatment, lost records, etc. A report by Prof. M. Thomas and Prof. H. Thimbleby [Ref 5] concluded that, had the attackers modified data rather than encrypting it, the attack could have gone unnoticed and caused significant harm.
- In 2014 researchers demonstrated how they could remotely access and control some elements of a standard Jeep. The extent of this control was limited, but the following year they further demonstrated how they could gain access to the control network that connects vital automotive systems including braking and steering [Ref 6]. As explained in Scientific American [Ref 7], these attacks required significant investment in time and skillset to accomplish, such that, whilst the effects could be quite damaging, care needs to be exercised to ensure a proportionate response.
- 2017 saw the first cyber attack specifically targeted at an industrial Safety Instrumented System (SIS). The 'HatMan'[38] malware [Ref 8] represented a sophisticated targeted attack that is estimated to have taken in excess of a year to develop and required knowledge of the target environment in depth. Whilst malware targeting Industrial Control Systems (ICSs) has been around since STUXNET [Ref 9], this is the first to be known to target an SIS. Although the HatMan malware triggered an automatic shutdown, there is speculation that the actual objective of the attack was to compromise the safety-related systems such that they failed to protect against an attack on control systems that would ultimately lead to a significant safety incident. In 2019 it was reported that an additional intrusion by the attacker behind the 2017 attack was reported at a different critical infrastructure facility[39].
- 2020 saw the emergence of a new type of ransomware in the EKANS[40] virus. The EKANS virus is a relatively straightforward piece of ransomware designed to encrypt files and display a ransom note. However, the EKANS virus showed additional functionality in which it searched for processes associated with ICS systems, and, if detected, the virus would forcibly stop the process. While this is a relatively primitive mechanism with a limited and fixed list of target ICS processes, it shows that the creators are aware of ICS systems as a target. Up until now, disruption caused to ICS systems from ransomware was more a result of collateral damage; however, EKANS represents a direct attempt to disrupt OT.

Attacks such as that using the HatMan malware demonstrate that some attackers are motivated to go to extraordinary lengths to compromise systems and are willing to cause harm in order to achieve their goals. The response from organizations needs to be proportionate to the risks, but there are many challenges in determining how to assess such risks and scale the response appropriately. This Code provides guidance on the consideration of proportionality.

Whilst HatMan represented a sophisticated and targeted attack requiring skilled resources, EKANS represents an evolution in threat and arguably requires a technical threshold for development and deployment. While the threat from EKANS in itself may initially be limited, it represents a development in that ICSs are now seen as targets by non-state developers of malware.

36 https://www.dailymail.co.uk/news/article-4503420/It-s-life-death-NHS-patients-say-cyber-attack.html
37 https://eandt.theiet.org/content/articles/2017/05/wannacry-and-ransomware-impact-on-patient-care-could-cause-fatalities/
38 Media reporting also refers to this malware as both TRITON and TRISIS.
39 https://www.fireeye.com/blog/threat-research/2019/04/triton-actor-ttp-profile-custom-attack-tools-detections.html
40 https://dragos.com/blog/industry-news/ekans-ransomware-and-ics-operations/

Annex B – Contemporary examples of threats and potential impact

Such attacks could be driven by a range of motivations, from extortion of the manufacturer, for example, through threatening damage to brand reputation, to theft, for example, obtaining a high value vehicle without having to expose the thief to physical law enforcement.

Examples of cyber security vulnerabilities that had a potential impact on safety are introduced below. Further details are available within the references cited:

- Announcements made in 2015 by the Cybersecurity and Infrastructure Security Agency (CISA)[41], part of the US Department of Homeland Security, raised concerns regarding the security vulnerabilities of a hospital drug pump, and the potential to alter drug dosing with lethal consequences. A security researcher highlighted further issues regarding other medical devices, including insulin pumps[42] and pacemakers. One manufacturer contacted customers to mitigate the potential threat of remote dosing from their insulin pump system. Reports have also highlighted the potential extraction of patient data from medical devices and their use to attack hospital networks.
- A 2016 incident reported by the US Food and Drug Administration (FDA)[43] responsible for medical devices described the life-threatening danger of security failures. A diagnostic computer monitoring, measuring and recording physiological patient data malfunctioned whilst being used for a cardiac catheterization procedure. There was a delay in the procedure whilst the application was rebooted. The FDA investigation found that communications between the patient device and the monitor were lost for five minutes while the patient was sedated, with no physiological data presented. Fortunately, the procedure was successfully completed after rebooting the application. However, the delay in care could potentially harm a patient. A configuration error of the anti-virus scan included directories that caused deletion of critical patient data.
- In 2017 a security consultant highlighted the potential for safety-related denial of service in industrial and medical surgical robots [Ref 33]. Researchers demonstrated eavesdropping and subsequent hijacking of communications ('Man in The Middle' attack) between the remote surgeon and robot in teleoperated robotic surgery. The researchers could take control and initiate an emergency-stop through fast movement (unsafe motion) or motion beyond zoned limits (safety areas). These actions caused the robot to shut down in a fail-safe mode, in the same manner as an industrial robot. This would force a reset of the safety system to commence further surgery. By sending malicious network traffic, the researchers were able to prevent the robot from being reset, preventing additional surgery from being performed. Automated devices with safety systems may have impacts as a consequence of security induced safety failures that have not been envisaged in traditional safety assessments.

Cyber attacks can use malicious software that exploits vulnerabilities in digital technology, often acting over digital networks. Table B.1[44] gives some examples of how security and safety objectives can be compromised by malicious software (for example, Virus, Worm, Trojan and Ransomware) being introduced onto the system. Note that this is just an example: many cyber attacks do not use malicious code – the adversary instead uses the system's own resources in unintended and unanticipated ways to attack the system.

41 https://www.us-cert.gov/ics/advisories/ICSA-15-125-01B
42 https://blog.rapid7.com/2016/10/04/r7-2016-07-multiple-vulnerabilities-in-animas-onetouch-ping-insulin-pump/
43 https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfmaude/detail.cfm?mdrfoi__id=5487204
44 Derived from NIST Cybersecurity Framework [Ref 10].

Annex B – Contemporary examples of threats and potential impact

Table B.1 Examples of the effect of malicious software

Event	Description/risks
Reduction in system availability and integrity	May reduce system resources available to a safety function or interfere with its operation.
Reduction in system availability and denial of control action	Device operation disrupted by intentionally or accidentally delaying or blocking the flow of information, denying device availability or networks used to control the device or system.
Reduction in system integrity, leading to configuration manipulation	Configuration settings modified e.g. to allow an attacker to gain increased access/ authority, producing unpredictable results or bypassing controls that mitigate operator errors by limiting their authority.
Reduction in system integrity, leading to spoofed system status	False information sent to operators either to disguise unauthorized changes or to manipulate operators into inappropriate actions.
Reduction in system integrity, leading to device functionality manipulation	Unauthorized changes made to embedded software, programmable instructions in devices, alarm thresholds changed, or unauthorized commands issued to devices, which could potentially result in damage to equipment (if tolerances are exceeded), premature shutdown of devices and functions, or even disabling equipment. Safety-related functionality manipulated such that they do not operate when needed; or perform incorrect control actions, potentially leading to physical harm or damage to equipment.

Annex C

Introduction to cyber security, safety and systems engineering

The following Sections provide a brief introduction to the topics that are central to this Code. Further detail is available by following the links included in Annex F, and in particular, Section F.2.3.

C.1 What is cyber security?

C.1.1 Defining cyber security

Cyber security is concerned with protecting digital systems and the value they generate. This can focus on the protection of information, technology and the human and organizational factors that depend on them. Given the ubiquity of digital systems in our modern society, cyber security affects many other concerns, including safety.

C.1.2 Protecting business value through protecting assets

The goal of cyber security is to protect an organization's objectives and its value. Ensuring that security activities support these high-level organizational concerns is necessary to validate their effectiveness and justify related expenditure. Any security risk management activity must flow from a clear understanding of the organizational objectives that are being protected, and the high-level impacts that the organization wishes to avoid.

In practice, cyber security is realised as a range of activities aimed at protecting specific organizational assets from cyber attack. These assets can include technology, information, money, people, policies, business processes, or anything else that is seen to generate value for the organization.

Where the asset in question is information, cyber security is concerned with protecting three key properties of that information[45]. These are confidentiality, integrity and availability.

C.1.3 Understanding cyber risk

When considering organizational assets, cyber risk is typically assessed according to three elements: threat, vulnerability and impact. Threat refers to the source of a cyber risk. This is typically a human threat, assessed according to their capability to mount an attack and their intent to do so. There are standardized taxonomies for describing cyber threats[46].

Vulnerability refers to a feature or flaw in an asset that may be exploited by a threat. This may take the form of a technical vulnerability, such as a website accepting code through a search box, which is then used to take control of the site's database. It could also take the form of a vulnerability in an organizational process or policy, such as the lack of a clear desk policy leading to sensitive documentary information being exposed.

Impact refers to the result of a threat exploiting a vulnerability. This can be described in terms of the impact on safety, finances, reputation, legal compliance or any other organizational concern.

These three concepts should be used to demonstrate how a cyber risk can impact upon the organization's goals. The purpose of this is to validate the necessity and sufficiency of risk reduction activities and to justify expenditure on those activities to reduce this risk to the organization's leadership.

45 Traditionally, these are confidentiality, integrity and availability. However, given the many different contexts in which digital technology is used in current society, these properties could be considered a restricted set of those required.

46 One example is STIX: https://oasis-open.github.io/cti-documentation/stix/intro

Annex C – Introduction to cyber security, safety and systems engineering

Cyber risk can be reduced by the application of security controls and by designing in a way that minimizes exposure to hazardous states, for example, by using resilient system engineering techniques like STPA. Many common cyber security standards include lists of controls known as control sets. These vary in their level of detail and the kinds of risk they aim to mitigate. Whilst many organizations operating in safety critical environments would seek to exceed the requirements of this Code, it is a good starting point for an organization when embarking on a programme of improving cyber security.

C.1.4 The management of cyber security

In order to be effective, cyber security needs clear management processes. These should be appropriate to the size and digital exposure of the organization in question. ISO 27001 provides one example of a template for a management system that can be used to control an organization's cyber security activities. These management systems typically follow continuous improvement lifecycles, such as 'plan, do, check, act', which can be aligned with other management systems, such as those applied to safety.

Within a Cyber Security Management System (CSMS), there is typically a defined set of activities referring to different stages of an organization's response to a cyber attack. These begin with preparatory work to identify cyber risks and to apply suitable controls to reduce those risks. In the context of specific attacks, activities are then defined to detect, respond to and recover from cyber attacks. It is noted that the application of these principles to operational technology is less mature than to traditional information technology.

C.2 What is safety?

Safety can be defined as the freedom from unacceptable risk of harm. Safety is traditionally considered in two main sub-disciplines: health and safety in the workplace and system safety. These two are not entirely unconnected, and the treatment of them will vary from organization to organization, dependent on the nature of the organization's activities. System safety can be further divided into functional aspects (hazards that arise from 'what it does') and passive aspects (hazards that arise from 'what it is'), although this latter differentiation becomes blurred when functions are used to provide controls for 'passive' safety issues.

NASA defines system safety as:

the application of engineering and management principles, criteria, and techniques to optimize safety within the constraints of operational effectiveness, time, and cost throughout all phases of the system lifecycle. System safety is to safety as systems engineering is to engineering. When performing appropriate analysis, the evaluation is performed holistically by tying into systems engineering practices and ensuring that system safety has an integrated system-level perspective.

It states that "The term 'system,' as used here, refers to one integrated entity that performs a specified function and includes hardware, software, human elements and consideration of the environment within which the system operates", emphasizing the holistic view.

Systems that have the potential to cause harm, or fail to protect from harm, occur in most sectors. These are referred to as safety-related systems[47]. The increasing use of software and complex digital systems presents an increasing challenge to understanding and assessing the safety risk. Traditionally, understanding the risk has involved modelling a causal chain from initiating event through to a hazard, and from that hazard to a postulated accident. Safety risk is expressed as a function of the severity of the harm of an accident and the likelihood of that harm being realised.

47 Some standards such as IEC 61508 [Ref 2] use additional criteria in their definition of a safety-related system. However, we use it here in its broadest natural language sense to be inclusive of any system that has an involvement in safety as a cause of, or as a barrier to, a harmful effect.

Annex C – Introduction to cyber security, safety and systems engineering

Safety risk reduction focuses on addressing the causal chain in a hierarchy of controls as shown in Table C.1.

Table C.1 Hierarchy of controls

Elimination	Redesign so that the hazard is removed or eliminated.
Substitution	Replace material or process with a less hazardous one.
Engineering controls	Design in barriers to the propagation of the causal chain. These can include a combination of: • Fault prevention • Fault tolerance • Hazard mitigation • Hazard recovery
Administrative controls	Identify and implement procedures you need to work safely.
Personal protective clothes and equipment	Protect the person at risk by using personal protective equipment (PPE).

Consideration must be given to all causes of hazards. Historically, it was considered sufficient to look at component failure, but now the nature of complex digital technology brings many other factors into consideration. For example, the ability of the operator to comprehend the current system state and make appropriate control demands to achieve their goal (or correct an anomalous situation) requires consideration of the human-machine interface (HMI). A study by the UK Health and Safety Executive (HSE) [Ref 19] concluded that the majority of root causes of incidents with safety-related control systems occurred in lifecycle phases where component failures were not directly relevant. This is illustrated in Figure C.1. Techniques have been introduced to address the complex socio-technical nature of modern systems. These aspects demand a holistic systems engineering approach that addresses the entire lifecycle.

Figure C.1 Primary cause by lifecycle phase
(Source: derived from HSE (*Out of control: Why control systems go wrong and how to prevent failure*; 2003)

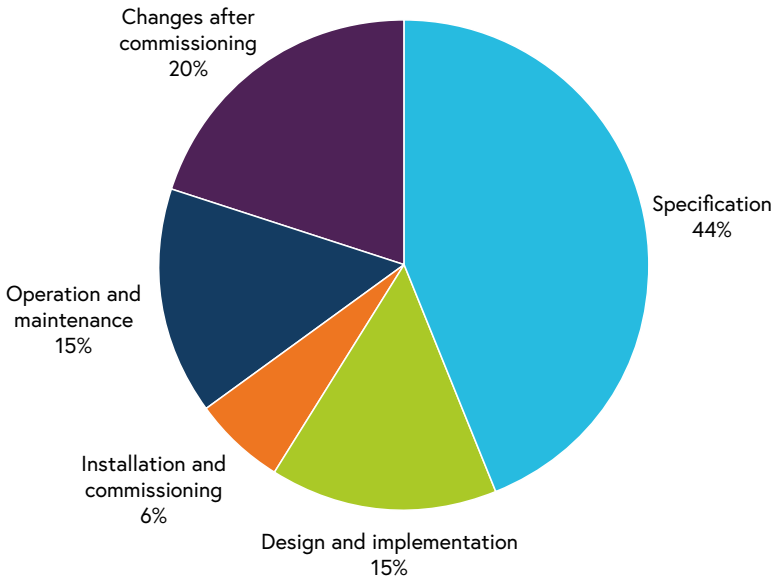

It is generally recognized that absolute freedom from risk of harm cannot be achieved, and so criteria are required to determine whether a safety risk can be accepted. This may be expressed in terms of the predicted likelihood of a particular level of harm occurring, and whether all reasonably practicable steps have been taken to reduce that risk. Often there is a need to trade risks, such as the risk benefit from introducing airbags to reduce the severity of an impact against the potential for harm caused by

the inadvertent inflation of an airbag. Absolute criteria are difficult to specify for the general case and safety standards will usually require a justification of the criteria used and their application.

This is further complicated by the challenges in assessing likelihood of effects where there are complex systems of hardware and people, and particularly where software is involved. The systematic nature of software makes it impractical to predict the likelihood of a software 'failure'[48]. Standards often specify controls on the software development and assurance processes on the premise that placing more rigour on these aspects will reduce the likelihood of a flaw being introduced in specification or implementation, and increase confidence that any such flaws will be detected and eliminated through the assurance process. Some standards also specify technical measures that increase resilience to unforeseen conditions and flaws that escape the development/assurance processes. Software safety standards are typically outcome-orientated in that they are not prescriptive of which combination of techniques and measures should be applied, and do not specify what likelihood of failure can be achieved by applying them. A justification of the selection and application of techniques and measures applied for a given design is required.

The non-prescriptive nature of legislation, regulation and standards has driven good practice to require documentation of the justification for the approach used to achieve a tolerable risk outcome. The justification needs to address all aspects, from selection of risk acceptance criteria, competence of those involved, suitability of development and assurance techniques and measures applied, design trade-offs considered, and through-life management. It is essential that bias in such justifications is addressed through independence and robust challenge throughout the lifecycle. It is useful to consider the justification in three areas [Ref 20]:

1 **Technical measures:** justification of measures taken to eliminate or manage hazards;
2 **Confidence:** justification of each step in the technical measure justification, e.g. considering completeness and adequacy; and
3 **Conformance:** justification of approach to compliance with legislation, regulation, standards, etc.

C.3 What is systems engineering?

The International Council on Systems Engineering (INCOSE) *Systems Engineering Handbook*[49] states that systems engineering is "a perspective, a process, and a profession".

Perspective	Systems engineering is an interdisciplinary approach and means to enable the realisation of successful systems. (INCOSE, 2004) • It focuses on defining customer needs and required functionality early in the development cycle, documenting requirements, and then proceeding with design synthesis and system validation while considering the complete problem: operations, cost and schedule, performance, training and support, test, manufacturing, and disposal. • Systems engineering integrates all the disciplines and specialty groups into a team effort forming a structured development process that proceeds from concept to production to operation. • Systems engineering considers both the business and the technical needs of all customers with the goal of providing a quality product that meets the user needs.
Process	Systems engineering is an iterative process of top-down synthesis, development and operation of a real-world system that satisfies, in a near optimal manner, the full range of requirements for the system. (Eisner, 2008)
Profession	Systems engineering is a discipline that concentrates on the design and application of the whole (system) as distinct from the parts. It involves looking at a problem in its entirety, taking into account all the facets and all the variables and relating the social to the technical aspect. (FAA, 2006)

48 Software does not fail in the sense that hardware does. However, it is convenient to talk about software failures where the behaviour of software does not match that expected or needed to achieve safety.

49 https://www.incose.org/products-and-publications/se-handbook

Annex C – Introduction to cyber security, safety and systems engineering

(Source: INCOSE *Systems Engineering Handbook*, 4th edition.)

Systems engineering addresses complex objectives by using abstraction to successively decompose from high-level concepts to address a desired capability through to concrete design implementations in hardware and software. Rasmussen [Ref 21] proposed an abstraction hierarchy that is helpful in understanding how we can relate the high-level purpose of a system to its real-world implementation. This is illustrated in Figure C.2. This is useful not just in the development decomposition, but also through in-service operation. The purpose of Rasmussen's original introduction of this hierarchy was to reason about how an operator in control of a complex system can make decisions in the event of anomalous system states. He proposes that complexity is a subjective feature of a system in that even simple objects become complex when viewed through a microscope. The objective of the application of the abstraction hierarchy is to achieve a system that is simple to operate to achieve the high-level purpose, even if the implementation technology is highly complex.

Figure C.2 Abstraction hierarchy (Source: derived from Rasmussen (The role of hierarchical knowledge representation in decision making and system management; 1970))

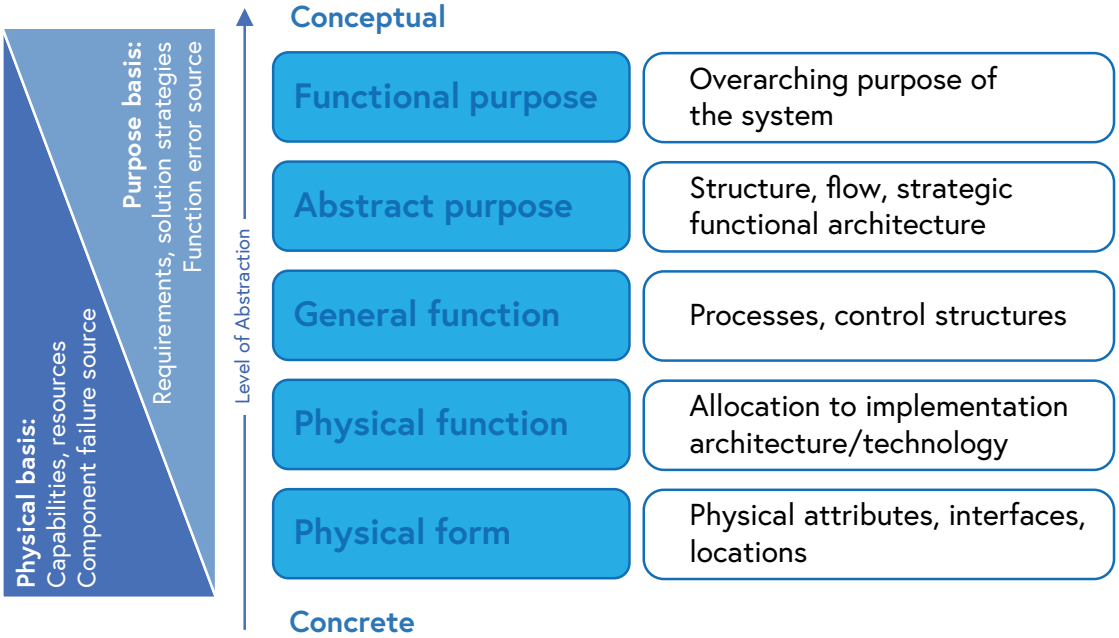

Levels of hierarchy are also used to help to manage responsibility for each step in this decomposition. Tiers of systems may be employed by an organization to achieve a desired capability within a business framework. ISO 15288 [Ref 15] uses a system concept where a system is composed of interacting system elements, in which a system element may itself be considered as a system, as illustrated in Figure C.3.

Annex C – Introduction to cyber security, safety and systems engineering

Figure C.3 ISO 15288 System-of-interest concept

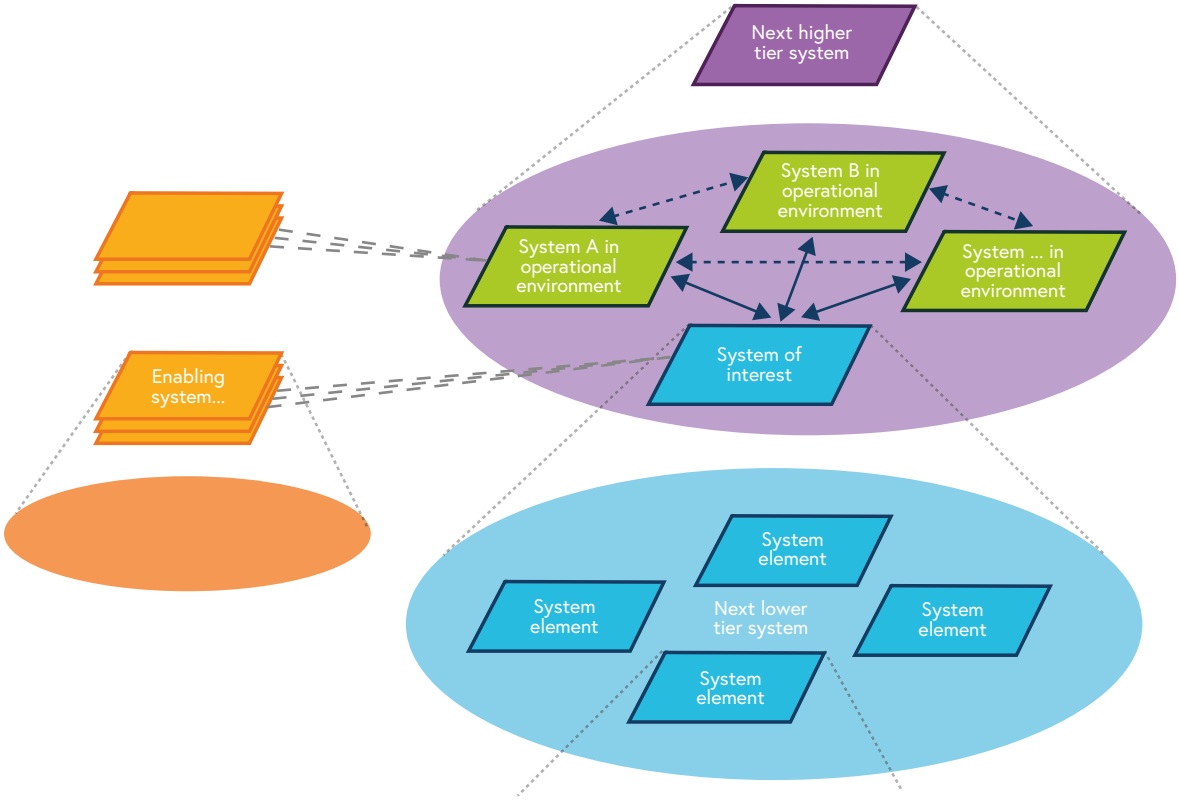

It adopts the term 'system-of-interest' to describe the current focus. A system-of-interest exists within an operational environment[50], which may comprise other systems, and is supported by essential services provided by enabling systems[51]. The system-of-interest may be decomposed into system elements that themselves comprise system elements or components. At the physical function/form level of abstraction, a system element can represent human activity such as that of an operator or maintainer. An organization decomposes the system-of-interest to the point where it identifies system elements that it can procure or build. These then become the system-of-interest for the supplier, with a supplier perspective on enabling systems, etc.

System-of-Systems[52] comprise independent constituent systems that serve their own purpose, but together produce a behaviour that cannot be achieved by individual systems. This adds a further level of complexity to dealing with safety and security.

50 In the context of this Code, the adversary exists within the operational environment, but may also have had malicious impact on the system-of-interest, its enabling systems or its sub-system/components.

51 It is noted that in the context of this Code, consideration must be given to the enabling systems as sources of threats and vulnerabilities, as well as considering the system-of-interest and systems in the operational environment.

52 For further information, see the INCOSE System of Systems Primer [Ref 22].

Annex D

Principles and indicators of good practice

The following table restates the principles and practices from Section 3 verbatim, and offers indicators of good and poor practice that may be used to make judgements about the need to change practices in an organization. It also relates the principles to the topics covered by the NCSC Cyber Assessment Framework (CAF). The indicators are set as examples and are not intended to be complete against the recommended practices or the CAF.

Annex D – Principles and indicators of good practice

Table D.1 Principles and indicators of good practice

The principle	The practice	Indicators of good practice	Indicators of the need to improve	CAF[53]
1. Accountability for safety and security of an organization's operations is held at board level.	**1** The board should put in place traceable delegation of responsibility and authority for addressing safety, security and their interaction. **2** The board should require regular and proactive reporting of issues that affect safety/security performance.	• One board member is identified with specific accountability for the complementary operation of cyber security and safety. • This role is widely recognized by staff.	• Different board members are accountable for cyber security and safety • There is no board awareness of the need for or discussion of the effectiveness of their complementary operation.	A1a
2. The organization's governance of safety, security and their interaction is defined.	**1** The board should set clear policies for safety and security. **2** The policies should encourage safety and security to be addressed co-operatively as part of an integrated systems engineering approach. **3** The organization should establish governance mechanisms to identify synergies and resolve conflicts between the objectives specific to safety and security.	• A policy document defines the responsibilities to enhance synergy and resolve conflicts between cyber security and safety. • Decision-making is prompt and effective.	• Middle and senior managers are unable to describe who is responsible for ensuring that cyber security supports safety. • Where the responsibility is assigned, posts are empty or the tasks are neglected.	A1b B1
3. Demonstrably effective management systems are in place.	**1** The organization should operate management systems that require the identification of relevant legislation and regulation. **2** The management systems should be designed to ensure that they identify inter-dependencies and interactions to ensure compatibility. **3** The Safety Management System and Security Management System should be parts of a comprehensive and coherent high-level management system. **4** The management systems should be maintained to ensure they manage safety and security risks using current relevant good practice. **5** The management systems should include measures to detect shortfalls against safety and security objectives and also identify and address the cause(s) of such shortfalls.	• There is a Cyber Security Management System (CSMS)[54] for the OT environment. • The SMS, CSMS and ISMS document their interdependencies. • The CSMS and SMS cite relevant legal and regulatory frameworks and describe their approaches to risk management.	• The SMS contains no reference to cyber security. • There is no CSMS. • The ISMS is assumed to cover the security of the OT environment. • Cyber security documents make no reference to supporting safety requirements and regulation.	A2
4. The level of independence in assurance is proportionate to the potential harm.	**1** The organization should establish criteria for the use of independent assurance against safety and security objectives, including the scope of independent assurance activities and level of independence.	• An independent assurance department considers systems aspects including safety and security together.	• Safety assurance and security assurance are assessed independently of each other.	A2

53 NCSC Cyber Assessment Framework (CAF) – https://www.ncsc.gov.uk/collection/caf

54 The term CSMS is used here to distinguish it from an ISMS that is more applicable to an information technology environment. The CSMS does not regard the Information as the primary asset to be protected.

Annex D – Principles and indicators of good practice

Table D.1 *Continued*

The principle	The practice	Indicators of good practice	Indicators of the need to improve	CAF
5. The organization promotes an open/learning culture whilst maintaining appropriate confidentiality.	**1** The organization should ensure that reporting of shortfalls against safety and/or security objectives in a trustworthy and responsible manner is encouraged. **2** The organization should ensure that the emphasis of investigation into shortfalls is seen as learning to avoid future shortfalls, rather than to allocate blame. **3** The organization should ensure a proactive approach to learning, through continuous improvement, training and sharing with the wider community.	• Management promotes learning from failures. • Documents demonstrate learning in action. • Security vulnerabilities are only publicized responsibly.	• Near-misses are hidden from management. • Those associated with bad news are punished. • Failures are obfuscated and the blame shifted.	D2 is closest
6. Organizations are demonstrably competent to undertake activities that are critical to achieving security and safety objectives.	**1** The organization should identify the key competencies required to achieve their safety and security objectives. **2** The organization should identify how the competency requirements are allocated across their organizational structure to groups and individuals with accountability, responsibility and/or authority for the setting or achievement of safety or security objectives. **3** The organization should record how it has ensured that the competency requirements are satisfied, and how these are maintained over time.	• Safety competencies include reference to cyber security and how it supports safety. • Cyber security competencies include reference to safety and how it is supported by security. • There are job-shadowing opportunities to ensure mutual understanding.	• Safety staff have no expectation placed on them to understand security. • Cyber security staff have no expectation placed on them to understand safety. • There is no routine interaction between safety and security personnel. • There is no mutual understanding between safety and security personnel.	B6 is closest
7. The organization manages its supply chain to support the assurance of safety and security in accordance with its overarching safety/security strategy.	**1** The organization should assess the nature of the relationship it needs with its suppliers in order to meet enduring obligations to provide cyber security services (e.g. patching, incident response support, etc.) for the lifetime of their products and services. **2** The organization should identify the requirements and responsibilities allocated to their supply chain to support achievement of the safety and security objectives. **3** The organization should identify the controls and reports it will use to manage the supply chain to ensure appropriate oversight of the factors that impact safety or security, including those that arise from emergent functionality/behaviour of the supplied product. **4** The organization should ensure it has sufficient in-house competence and capacity in any outsourced safety or cyber security assessment and advice services that it can retain control of its risk decision-making.	• The CSMS and SMS explain how the organization's approach to complementary security and safety extends to the supply chain. • Suppliers can describe the organization's approach to and processes for complementary safety and security.	• Safety requirements are found in one contract annex, cyber security requirements are in a different annex and there is no evidence of coherence. • Contracts can demand conflicting outcomes from safety and security and the organization is unaware of these conflicts. • Risk management decisions are made on the advice of out-sourced advisors without comprehension or challenge.	A4

Annex D – Principles and indicators of good practice

Table D.1 *Continued*

The principle	The practice	Indicators of good practice	Indicators of the need to improve	CAF
8. The scope of the system-of-interest, including its boundary and interfaces, is defined.	**1** The organization should identify, document and communicate the scope of the system-of-interest within the bounds of its safety and security objectives. **2** The organization should identify interacting systems in the operational environment and the enabling systems that could impact safety or security objectives.	• The system-of-interest that requires both safety and security to be assured is identified in technical and organizational terms. • The system boundary is enforced through technical, procedural and organizational means. • The interfaces and dependencies across that boundary are defined, understood and managed.	• The boundary of the system-of-interest is unclear to staff. • Ownership and organizational responsibilities at the boundary are unclear, e.g. for boundary firewall rules. • Dependencies, e.g. from IT into OT for mission-critical network services, are not documented with Service Level Agreements (SLAs).	B4
9. Safety and security are addressed as co-ordinated views of the integrated systems engineering process.	**1** The organization should define engineering and business processes that enable separate security and safety disciplines to co-ordinate their activities against the safety and security objectives.	• The SMS and CSMS share common documents defining and describing the system-of-interest. • The safety and security lifecycles indicate interdependencies in their respective management systems.	• There is no common understanding of the system-of-interest and any safety-security discussion starts with a debate about what they are talking about. • There is no routine interaction and no recognized interdependencies between safety and security.	B5 is closest
10. The resources expended in safety and security risk management, and the required integrity and resilience characteristics, are proportionate to the potential harm.	**1** The organization's safety and security management systems should define frameworks and criteria that guide engineering activities to achieve a proportionate approach.	• Safety and security staff recognize the need to work together to determine how to assess the proportionality and sufficiency of security measures to support safety requirements. • Safety and security staff understand the principles and basis of each other's analysis.	• Safety staff assume that the performance of security measures can be specified and measured as for safety, e.g. with Probability of Failure on Demand figures. • Security staff are unaware of the nature of safety analysis and that the adequacy of measures may be tested against 'reasonable practicability' criteria.	B5

Table D.1 *Continued*

The principle	The practice	Indicators of good practice	Indicators of the need to improve	CAF
11. Safety and security assessments are used to inform each other and provide a coherent solution.	**1** Techniques should be selected that are appropriate to the lifecycle point and the objective of the assessment. **2** Interaction points between security and safety assessments should be identified and communicated. **3** The results of assessment, whether quantitative or qualitative, should be expressed together with a measure of confidence.	• The outputs of process hazard analysis procedures, such as a Hazard and Operability study (HAZOPS), inform cyber security analysis. • HAZOPS can consider multiple contingencies to cover malicious action. • The generally and inevitably more qualitative nature of security risk assessment is recognized.	• There is no interaction of safety and security in the identification of risks to OT. • The safety risk analysis does not consider risks arising from malicious action, e.g. as multiple, coordinated initiating events. • Security analysis does not consider the outputs of hazard analysis activities and focuses entirely on hardening systems and networks.	A2
12. The risks associated with the system-of-interest are identified by considerations including safety and security.	**1** Risk identification should be performed using team-based techniques employing a broad range of competencies and viewpoints. **2** Techniques that provide a systematic basis for risk identification should be employed. **3** All identified undesirable outcomes, together with assumptions or identified interactions, should be recorded before being rationalized. **4** Risks should be rationalized to remove duplication and produce a coherent set of risks to be addressed. **5** The process of rationalizing risks should identify system hazards at a consistent level of abstraction. **6** Care should be taken not to rationalize risks on the basis of lack of a credible cause, or the extreme unlikelihood of the cause occurring.	• The identification of risks involves collaboration between disciplines. • Risk identification techniques are chosen explicitly to lead to a common understanding between safety and security of what system states are to be avoided. • There is a company policy strictly limiting the use of likelihood to discount risks, in accordance with the relevant regulations and guidance, e.g. for COMAH sites in the UK.	• There is no consideration of achieving a common understanding between safety and security of hazardous system states. • Assessment of security risks for specific OT is largely agnostic of the business functions controlled by the OT. • Risks are often excluded from further analysis because nobody can imagine the sequence of events to cause the hazardous state. • New technology is adopted without adequate understanding of its safety/security implications.	A2

Annex D – Principles and indicators of good practice

Table D.1 *Continued*

The principle	The practice	Indicators of good practice	Indicators of the need to improve	CAF
13. System architectures are resilient to faults and attack.	**1** System architectures should be engineered to employ multiple layers of protection that make attacks more difficult, more likely to be detected and responded to and less likely to lead to harm. **2** Maintenance and operations activities should have an enduring obligation to ensure that the engineered safety and security protections are not compromised or circumvented. **3** Systems and components should be designed such that security controls can be maintained and updated, in operation, in a safe manner. **4** Components should be hardened against targeted attacks, according to their role in fulfilling the safety and security objectives. **5** The possibility that mitigation and recovery systems may be affected by a security attack should also be considered. **6** Organizations should clearly identify and communicate the events and activities that lead to a decision to move the safety system into or out of reduced functionality mode (to protect the integrity of the safety system against attack).	• Safety Layer of Protection Analysis (LOPA) informs and is informed by cyber security analysis. • The identification of cyber security measures, e.g. zoning and strength of measures, is influenced by safety criticality, e.g. by function Safety Integrity Level (SIL). • Heightened security states and reduced functionality safety states are designed in a co-ordinated way. • Maintenance and Operations understand the complementary nature of safety and security.	• Independent safety protection layers may be found to have unrecognized common mode security weaknesses. • There is no recognition in security risk analysis of the need to protect safety-critical systems except according to the sensitive information they contain. • Safety response and security response are entirely independent. Nobody knows what will happen when they coincide.	B5, D1
14. The risk justification demonstrates that the safety and security risks have been reduced to an acceptable level.	**1** Risk criteria should be established by the organization that bound tolerable levels of risk against safety and security objectives. **2** The residual risk of harm against safety and security objectives should be justified against the risk criteria. **3** A holistic approach should be taken to the justification of system trade-offs, seeking alternatives that provide optimal satisfaction of all required properties, including achievement of safety and security objectives.	• The company has policy and guidance on risk criteria[55] for cyber threats to safety. • Active steps are taken to identify and resolve tensions between security and safety controls – and with implementation – as early as possible. • Synergies are exploited, e.g. security cameras for safety.	• There is no argument to justify why security measures are necessary and sufficient to support safety. • Unrecognized conflicts between safety and security controls exist in the implementation, e.g. between fail-safe and fail-secure.	A2
15. The safety and security considerations are applied and maintained throughout the life of the system.	**1** Activities to address safety and security objectives should be planned and addressed at every stage of the entire lifecycle, from concept to disposal. **2** Activities to address safety and security objectives should be addressed for every aspect from architecture to component level, and from concept to disposal. **3** Planning for risk assessment review should recognize that the evolution of the threat environment is potentially more dynamic than that of the technology and operating environment. **4** Change management should be an active process operating from concept to disposal.	• Safety staff understand how security engages in the safety lifecycle. • Similarly for security. • Design, operations and maintenance similarly, with their own lifecycles. • Safety assumptions are reassessed when supporting security assumptions change.	• Safety, security, system lifecycles do not have the necessary interactions to achieve complementary safety and security. • Discovery of new security vulnerabilities and threat scenarios trigger a reassessment of security measures only.	B6 B1

55 In the UK this requires a working equivalent to ALARP for security.

Annex E

Techniques and measures

This Annex contains introductions to techniques and measures that may be used to help assess, understand or manage the risks arising from where security impacts on safety.

It addresses the following topics:

E.1 Risk control systems

E.1.1 Introduction

Type: risk management framework

A 'risk control' is the method by which an organization evaluates potential losses and harm, and takes action to eliminate, reduce and control vulnerabilities and hazards throughout a system, service or equipment life. The hierarchy of cyber and safety risk controls, starting at the most effective, are:

(a) elimination;
(b) substitution;
(c) passive engineering;
(d) active engineering; and
(e) operational.

Regulators, government departments, institutes and organizations promote risk reduction and control approaches for both cyber and safety vulnerabilities and hazards, respectively. The risk controls for both cyber and safety are on the whole coincident. For most organizations that work across multiple domains and need to satisfy multiple regulations, it is worth addressing the risk control systems identified below as part of standard business practices. It is likely that most organizations will already have some of these risk control systems in place. However, more often than not, there are separate systems for the safety and cyber functions.

Where safety-involved systems including industrial control and automation systems are in scope of the essential service, controls suitable for managing risks on the corporate information technology network may be inappropriate or damaging in an operational technology environment. These systems require a more tailored approach to address the potential hazards.

E.1.2 When to apply it

As a risk management framework, consideration of risk control should be applied throughout the lifecycle from concept to disposal.

E.1.3 Basic method

Organizations should have a systematic process in place to ensure that identified risks are managed and the organization has confidence that mitigations are working effectively. An NCSC principle is that organizations take appropriate steps to identify, assess and understand security risks to the network and to information systems supporting the delivery of essential services. This includes an overall organizational approach to risk management. It is recommended that organizations embed cyber safety risk control systems in their business management and systems engineering processes and procedures. Good practice can be found in the links available in Further information (Section E.1.6 below).

E.1.4 How it can be adapted

Risk control systems	Definition
Governance and management systems	This defines the senior management commitment, leadership and risk ownership. It also provides the organization's strategic direction, oversight and decision-making, setting out core policies that define an organization's purpose, values and structure. The management systems enable the strategic direction and provide the governance organization. As part of an organization's risk management and governance processes, cyber security and safety need to be continually assessed, particularly as the threat environment continually changes, as will investment priorities. This includes the identification of dutyholders that have legal duties under The Network and Information Systems (NIS) Regulations 2018 as Operators of Essential Services (OESs).
Regulatory approvals and certification	Regulators such as the HSE and the ONR require the consideration of cyber security as a component of safety risk and organizations are required to demonstrate that the safety risks that could result from cyber vulnerabilities have been addressed.
System, equipment, service upgrade and maintenance plans	Safety-involved OT is likely to have followed a functional safety approach such as IEC 61508 *Functional safety of electrical/electronic/programmable electronic safety-related systems*, or one of its industry/application-specific variants such as IEC/EN 62061 *Safety of machinery*. These Standards require quality control, management processes, validation and verification techniques, and failure analysis, etc. processes at a level of rigour that depends on the derived integrity levels. The design of the systems needs to provide for the management of cyber security risks through the updates that fix emergent security vulnerabilities, so that at-risk times are minimized and the safety integrity remains assured. A close relationship is required with the equipment supply chain to ensure vulnerabilities are understood and addressed in a timely and safe way. This starts with understanding all the safety-related Operational technology and information technology assets and their cyber vulnerability status.
Roles and responsibilities	The dutyholder needs to ensure that adequate competent and empowered resources are available to manage cyber security risks, including effective management systems and technical countermeasures.
Organizational competency	The ability of an organization to support its cyber and safety people through its governance, structure, control and investment systems, knowledge management and good practice processes and procedures, providing a learning culture captured in its business management systems. This may be recognized in part through approvals such as the ISO 9001 quality management system, ISO 27001 *Information security management*, ISO 55000 *Asset management* and Cyber Essentials certification.
Competency and training	The ability to undertake responsibilities and to perform activities with regard to specific standards to meet the task. See Annex E.2 for more information.

Annex E – Techniques and measures

Risk control systems	Definition
Requirements management	The hazard log is a key output once the application of the hazard management process is undertaken, as it should provide evidence that safety requirements have been met, and hence that the hazards have been closed. It can be modified to include the related vulnerabilities and resulting mitigations that should also be translated into cyber requirements. Requirements may include any of the following: • implementation features of functions of technical systems, including defensive architectures and segregation; • delivery of minimum levels of integrity (assurance) for functions of technical systems and how this will be maintained during patching and updates to address threats; • the operational arrangements, including organizational, such as countermeasures, access controls, identity, data security, provision of user manuals, zoning, back-up, provision of training, updates to operational procedures and limitations on use; • the maintenance arrangements, including the provision of tools, patch updates, spares, special equipment, maintainer training, and inclusion of certain checks within maintenance procedures; and • any restrictions introduced for system operation to control risk while assumptions about the behaviour of the system are confirmed. Compliance with a requirement may be demonstrated by trials, testing, inspection or analysis.
Risk identification and management	The cyber threat changes over time and it can be difficult to define the potential safety consequence for each threat. A continuing detailed analysis is required as threats change, noting that the likelihood of new threats arising is unpredictable. Such analysis can use safety risk techniques, such as Fault Tree Analysis, Event Tree Analysis and Bowtie diagrams, to inform cyber security risk techniques, such as Attack Trees and cyber vulnerability analysis. However, normal risk assessment processes such as hazard and operability studies or process hazard risk analysis, etc. are not likely to be sufficient to address cyber security threats, as these approaches do not always consider several dangerous events occurring at once or address those that have malicious intent. In addition, the loss of essential services with respect to the impact on national infrastructure needs to be considered. Techniques such as System-Theoretic Process Analysis (STPA) may provide a more holistic view of the potential hazardous events.
Recognized design and manufacture good practice	Appropriate national or international codes and standards should be adopted for safety, security, systems and equipment. Where necessary, the organization should develop its own engineering standards and supporting guidelines to adapt these codes and standards to their business. The codes and standards applied should reflect the principle that the functional reliability of systems and equipment should be proportionate with their security risk. As part of any ALARP argument, the relevant good design, manufacturing and operational practice should be followed as a minimum. In many cases it is likely that relevant good practice (accepted by the regulator) as reducing risks to ALARP will have already been established. The effects of the cyber vulnerabilities on the good practice should be assessed and countermeasures put in place to maintain the safety argument.
Secure by Design	'Secure by Design' is an approach that seeks to reduce vulnerabilities during the definition and design process rather than trying to secure the system after setting it to work. It aims to mitigate specific threats by using a hierarchy of controls approach, with design or arrangement tailored to address malicious acts. It includes: • understanding the existing infrastructure; • security layers; • segregation of safety-involved systems; • designed-in privilege enforcement; • equipment connectivity; • design for maintainability; • third party risk management; • human factors, training and awareness; and • protected safe operating limits.
Control of hazardous materials	Many safety regulations require that hazardous materials are adequately controlled. Clearly, protecting this material to maintain a safe state, including securing its data, is paramount.

Annex E – Techniques and measures

Risk control systems	Definition
Product/system boundary and control responsibilities	Clearly defining the extent of the dutyholder's responsibilities and those of the interfacing organizations is critical to ensuring that risk ownership is understood. Identification of any reliance on other stakeholders, particularly for common cause risk (e.g. electrical power) is key to ensuring that appropriate mitigations are in place.
Change management	The change management procedures should control any additions, deletions or modifications to the safety-involved cyber systems and equipment. The objective is to ensure that standardized methods and procedures are used for efficient and prompt handling of all changes in order to minimize the number and impact of any related cyber and safety incidents upon the business.
Maintenance of condition	We can become 'normalized' or accustomed to seeing situations when safety and cyber security systems are damaged, degraded or not maintained. Tolerating small defects can lead to major incidents over time. Regular formal reviews and events are required to reinforce the continued maintenance of condition.
Tests and trials	Tests and trials may be used to demonstrate that the requirements are effective in mitigating the hazards and the vulnerabilities. In addition to the safety tests, these include cyber security tests and evaluation activities, including vulnerability assessments, security controls testing, penetration testing, adversarial testing, cyber security testing related to a system's resiliency, emergency preparedness and survivability capabilities.
Documentation, data and evidence	Information security, particularly for safety information and associated assets (including equipment, operational technology, information technology and data), needs to be managed. It is important that documents, data and supporting evidence are classified and protected using a defined policy. Information that is stored, processed or transmitted by digital systems needs to be made secure. The risks posed by portable media devices need to be identified and appropriately managed.
Safe and secure operating envelope	Under normal operating conditions, the system and equipment should remain within its safety envelope. When the system state trend predicts that it will fall outside the safety limits, the monitoring and detection should be engineered to remain in a safe state. Also, operational processes should be in place to manage the continued safe operation or shut down of the system.
Reporting and trend analysis	To enable good governance, data is required about the effectiveness of the risk control systems. It is important to collect data on the risk control systems and then decide what data should be analysed. The time frame for the analysis should be selected; for example, this could be continually to yearly. Data visualization can be plotted against processes, systems and equipment to aid understanding. Care needs to be taken when applying advanced trend analysis techniques, as these can be prone to error and false deductions – correlation does not always equal causation.
Emergency preparedness	Emergency Preparedness and Response (EP&R) planning should be in place to take all reasonably practicable measures to prepare for possible security emergencies and to mitigate the consequences.
Assurance Case – Snapshot of risk control effectiveness	A structured argument, supported by a body of evidence that provides a compelling, comprehensible and valid case that a system is safe and secure for a given application in a given operating environment. Contingency planning should address the immediate response, consequence management, maintaining situational awareness, effective command, control and communications.
Quality assurance and audit	The cyber security/safety processes should be integrated into the organization's business management systems and be subject to audit and certification in accordance with ISO 9001 *Quality management systems* and ISO 27001 *Information security management systems*.
Independent verification and validation	For safety-involved systems, the use of Independent Safety Auditors is well understood and defined as a role. The same value and support to the cyber security/safety argument can be obtained by the use of an Independent Cyber Assessment (ICA). An ICA is the formation of a judgement by a recognized subject matter expert, separate and independent from any system acquisition, design, development or operational personnel, that the requirements for the system are appropriate and adequate for the planned application and that the system satisfies those requirements. In discharging this responsibility, the key tasks for the function are: • acquiring an appreciation of the scope and context of the assessment; • selecting and planning a cost-effective assessment strategy; • gathering relevant evidence; and • forming a judgement, including managing any outcomes.

Risk control systems	Definition
Culture	Culture is defined as shared values and beliefs that interact with an organization's structure and control systems to produce behavioural standards. It starts with senior management acting as a living embodiment of the culture and leading by example. Having a positive safety and cyber security culture is key to managing the risks. A 'just culture' enables the organization to have the best interaction with its people about safety and cyber security. People are encouraged to speak up and report concerns and appropriate action is taken.
Regulatory approvals and certification	Regulators such as the HSE and ONR have put in place principles for dutyholders and guidance for their inspectors. It is expected that, as this domain matures, safety-involved systems will form part of the approval, certification and enforcement regime.

E.1.5 How it helps

Shared understanding of risk control systems and their terminology will aid communication between the safety and security communities.

E.1.6 Further information

- HSE OG-0086 *Cyber Security for Industrial Automation and Control Systems* (IACS), Edition 2 [https://www.hse.gov.uk/foi/internalops/og/og-0086.pdf]
- ISO 27001 *Information security management systems*. This provides good practice to address information security that encompasses people, processes and technology that should result in an effective Information Security Management System (ISMS).
- *Cyber Assessment Framework (CAF) guidance*: National Cyber Security Centre (NCSC) [https://www.ncsc.gov.uk/collection/caf]
- *10 steps to cyber security*: NCSC. This guidance helps organizations protect themselves in cyberspace. It breaks down the task of defending networks, systems and information into its essential components, providing advice on how to achieve the best possible security in each of these areas. [https://www.ncsc.gov.uk/collection/10-steps-to-cyber-security]
- HSG65:2013 *Managing for health and safety*. This guidance explains the 'Plan, Do, Check, Act' approach and shows how it can help you achieve a balance between the systems and behavioural aspects of management. It also treats health and safety management as an integral part of good management generally, rather than as a stand-alone system. [https://www.hse.gov.uk/pubns/books/HSG65.htm]
- ISO 45001:2018 *Occupational health and safety management systems – Requirements with guidance for use*
- RTCA DO-355 *Information Security Guidance for Continuing Airworthiness* covers operations and maintenance
- RTCA DO-326A/EUROCAE ED-202A *Airworthiness Security Process Specification*
- RTCA DO-356A/EUROCAE ED-203A *Airworthiness Security Methods and Considerations*
- ISA-TR84.00.09-2013 *Security Countermeasures Related to Safety Instrumented Systems* (SIS)
- BS EN/IEC 62443 *Security for industrial automation and control systems*
- *Cybersecurity Test and Evaluation Guidebook* (Version 2.0): US Department of Defense [https://www.dau.edu/cop/test/_layouts/15/WopiFrame.aspx?sourcedoc=/cop/test/DAU Sponsored Documents/Cybersecurity-Test-and-Evaluation-Guidebook-Version2-change-1.pdf]
- *Cybersecurity Test and Evaluation Guidebook* (Overview presentation 2018): US Department of Defense [https://www.dau.edu/cop/test/DAU Sponsored Documents/Cybersecurity Standard Overview 25 July2018 Marked.pdf]
- *Promoting a positive culture – A guide to health and safety culture*: Institution of Occupational Safety and Health (IOSH) [https://www.iosh.co.uk/~/media/Documents/Promoting a positive cultureconnect.pdf]
- *Board Toolkit – Developing a positive cyber security culture*: NCSC [https://www.ncsc.gov.uk/collection/board-toolkit]
- *Security Assessment Principles for the Civil Nuclear Industry* (Version 0): Office for Nuclear Regulation (ONR) [http://www.onr.org.uk/syaps/index.htm]

E.2 Competencies

E.2.1 Introduction

Type: governance; indirect confidence building enabler

In any well-run organization, staff are required to be competent to perform the tasks assigned to them. This is particularly important in organizations dealing with cyber risk that could result in harm where their regulators and customers need assurance that the organization's personnel can be shown to meet the necessary standards of competency.

There are several benefits to assessing formally the competence of individuals, including meeting the requirements of legislation and regulation, identifying personnel development needs, balancing team skills and as an aid to succession planning.

Competence is not just about qualifications. It also includes skills (for example, problem-solving) and behaviours (for example, personal integrity) that enable an individual to perform a function effectively. The following Sections provide guidance on developing individual/team competencies rather than the organization's competencies.

E.2.2 When to apply it

Consideration of competencies for key people and teams should be applied throughout the lifecycle, from concept to disposal. It should be recognized that differing levels of competence in cyber security and safety are required in non-engineering roles such as Operations Managers, Maintenance Managers, etc. as well as within engineering and operator roles.

E.2.3 Basic method

It is recommended that organizations develop specific competencies by combining industry recognized good practice. This good practice includes:

- Chartered Institute of Information Security (CIISec) Skills Framework
- NCSC Certified Professional (CCP) scheme and Certified Cyber Security Consultancy (CCSC)
- IET *Code of Practice: Competence for Safety-Related Systems Practitioners* (2016)
- Skills Framework for the Information Age (SFIA)
- HSE: *Managing competence for safety-related systems. Part 2: Supplementary material*
- BCS: *A Pragmatic Guide to Competency: Tools, Frameworks and Assessment* (Jon Holt and Simon A. Perry, 2011)

The organization should set out the competencies expected and the evidence required to assure competence in specific tasks. It should also create schemes for monitoring and measuring the competencies of its employees.

E.2.4 How it can be adapted

The roles that the organization requires to undertake the necessary tasks should be identified and the reporting lines defined. The definition of these roles will aid the identification of the skill sets, including any basic educational needs. Typically, a role is defined as follows:

Purpose: Cyber Security Engineering ensures that information technology and connected operating technology systems (including hardware, software, operators, maintainers and data) are hardened and resilient to cyber attacks. Cyber Security Engineering is a sociotechnical discipline that:

(a) Ensures that the development, operation and disposal of products, services, platforms or systems and supporting facilities are hardened and resilient to cyber attacks;

(b) is informed by the analysis of threat, vulnerability and business impact;
(c) applies risk control systems, including relevant good cyber resilience practice;
(d) is founded on a knowledge of attack modes that can contribute to loss and harm; and
(e) supports through-life cyber risk management.

This is then expanded to include **responsibilities** in terms of the tasks and outputs used to enable good decisions. This practitioner role focus facilitates self-management of the practitioner's professional competencies, tempered by being appointed by a peer, and is a pivotal Level. Table E.1 provides a typical definition of a practitioner role for illustrative purposes:

Table E.1 Illustrative practitioner roles for competence management

Item	Responsibility	Outputs
1	Understands the business, its operational environment, its information systems, equipment and data and the potential vulnerabilities.	Cyber security/safety strategy
2	Is the technical authority within the business for cyber security/safety and is able to plan for the implementation of the cyber security/safety strategy and supports governance activities.	Inputs into the: • security plan • through-life engineering management plan
3	Produces models and architectural representations that allow stakeholders to identify the potential threats, vulnerabilities and business impacts.	• Sociotechnical systems models and architectures • Checklists
4	Undertakes vulnerability investigation and analysis with stakeholders using techniques such as STAMP, SWIFT, Attack Tree, Data Flow Analysis, HAZOPS, etc.	• Cyber security/safety system vulnerability analysis • Cyber security/safety vulnerability investigation report
5	Defines industry-recognized good cyber resilience practice, standards, techniques and architectures to mitigate cyber safety risks, including zoning, Identity and Access Control, Data Security, Network System Security and Security Monitoring.	Inputs into the: • design definition • operational documents • maintenance procedures
6	Leads the development of all required system cyber documents and collates them in a System Security Case to present an overall risk-based argument for the fulfilment of the equipment cyber resilience requirement.	Input into the security and safety case(s)
7	Assesses system changes (technology or process) for potential cyber safety risks and proposes effective mitigations.	Cyber security/safety system vulnerability analysis
8	Manages through-life cyber risk management.	System accreditation, certification and approval
9	Audits and assesses cyber security/safety activities through life.	Audit report
10	Defines the organization's processes, tools, procedures and training.	• Processes • Tools • Procedures • Training
11	Carries out incident investigation and makes recommendations to mitigate vulnerabilities.	• Incident reporting

Competency levels: in order to identify coherent competencies, the number of competency levels must be determined. Typically, the levels of competence can be characterized as:

- ***Aware*** – an essential level of skill to understand the importance of cyber security to the organization. Competent to support threats, vulnerabilities and business impact analysis. Requires minimum training aimed at awareness not understanding.
- ***Supervised practitioner*** – an aspirant practitioner who therefore needs the same training. Guided and mentored to the standard required. Uses a skills management system to record progression and competency achievement.
- ***Practitioner*** – self-management of professional competence tempered by being appointed by peer review. Uses a skills management system to record progression and competency achievement. Provides on the job training and knowledge transfer to supervised practitioners.

- ***Expert*** – recognized outside the organization; training for continuing professional development routinely only cost-effectively delivered by external organizations, due to the small number of people. Uses a skills management system to record progression and competency achievement. Provides on the job training and knowledge transfer to practitioners.

The following skill groups are likely to be used to define the **competencies**:

- governance;
- policy and standards;
- equipment cyber resilience strategy;
- innovation and business improvement;
- equipment cyber resilience awareness and training;
- legal and regulatory environment;
- supply chain management;
- risk assessment;
- risk management;
- architecture and modelling;
- resilient by design;
- resilient systems and equipment development;
- assurance methodologies;
- testing;
- safe and secure operations management;
- safe and secure operations and service delivery;
- vulnerability assessment;
- incident investigation and management;
- forensics;
- audit and review;
- business continuity planning; and
- research.

E.2.5 How it helps

It helps to ensure that organizations have the right people in place to manage the safety-cyber risk to exceed the minimum standards set by legislation and regulation.

E.2.6 Further information

- *A Pragmatic Guide to Competency: Tools, Frameworks and Assessment* [https://www.amazon.co.uk/Pragmatic-Guide-Competency-Frameworks-Assessment/dp/1906124701]
- *Managing competence for safety-related systems. Part 2: Supplementary material*: HSE [https://www.hse.gov.uk/humanfactors/topics/mancomppt2.pdf]
- *Code of Practice: Competence for Safety-Related Systems Practitioners*: IET [https://shop.theiet.org/code-of-practice-competence-for-safety-related-systems-practitioners]
- *A guide to creating your own competence framework*: IET [https://www.theiet.org/media/1804/a-guide-to-creating-your-own-competence-framework.pdf]
- SFIA [https://www.sfia-online.org/en]
- BCS SFIAplus – IT skills framework [https://www.bcs.org/develop-your-people/develop-your-team-or-organisation/sfiaplus-it-skills-framework/]
- Chartered Institute of Information Security (CIISec) Skills Framework [https://www.ciisec.org/]
- HSE Research report 086 *Competence assessment for the hazardous industries* [https://www.hse.gov.uk/research/rrpdf/rr086.pdf]
- Certified Professional scheme [https://www.ncsc.gov.uk/information/about-certified-professional-scheme]
- Certified Cyber Security Consultancy [https://www.ncsc.gov.uk/information/ncsc-certified-cyber-security-consultancy]

E.3 System-Theoretic Process Analysis (STPA)

E.3.1 Introduction

Type: qualitative hazard analysis technique; applies system theory to complex interactions of system components

System-Theoretic Process Analysis (STPA) is a well-established technique developed by Professor Nancy Leveson and colleagues at Massachusetts Institute of Technology (MIT) for assessing safety and security risk. It can provide a common holistic method around which the members of the multi-competency team can unite to reason about safety and security and consider the effects of such things as the loss of equipment, erroneous operation, disruptive information flows, deliberate unsafe human interaction, attack vectors and how system constraints could be used to help ensure continuing safe and secure operations.

E.3.2 When to apply it

STPA should be applied throughout the lifecycle from concept to disposal to assist in identifying safety requirements and constraints. STPA can be used to help make design decisions from high levels of abstraction in concept, through architectural decomposition to detailed design. STPA includes software and human operators in the analysis, ensuring that the hazard analysis includes all potential causal factors in losses and is readily integrated into system engineering processes.

E.3.3 Basic method

The STPA method builds on four simple steps:

1. Define the purpose of the analysis;
2. Model the control structure;
3. Identify Unsafe Control Actions (UCAs); and
4. Identify loss scenarios.

E.3.3.1 Define the purpose of the analysis

Defining the purpose starts with an understanding of the system boundary and the environment in which it operates. The boundary is set such that everything inside the system boundary is under the system developer's control, whilst they are assumed to have no direct control over the environment.

Losses are identified in terms of outcomes that are of value to the stakeholders. They should not reference system elements or causes. It is the goal of STPA to prevent losses. The choice of language is deliberate, to ensure it is agnostic of any particular discipline. Whilst in safety analysis we might consider loss of life or injury to people, losses can equally be used to describe asset damage/loss, mission loss, etc. Where more than one loss is identified, they can be ranked and prioritized.

From the losses, system-level hazards are identified. These represent the system state that, together with a particular set of environmental conditions, will lead to a loss. Hazards are expressed in terms of the system states to be avoided. They should not include causes or failures, or include ambiguous wording such as 'unsafe'. The goal of safety engineering is to control the potential for each hazard and/or their effect.

System-level constraints are identified such that when they are enforced, they will prevent or control each hazard, and such that the collection of all system-level constraints are sufficient to provide adequate control of all system-level hazards. Constraints are solution-agnostic and define what needs to be achieved rather than how it is achieved.

In large/complex applications, system-level hazards can be refined into 'sub-hazards' by consideration of the basic system processes/activities. This allows more manageable hazards and constraints to be captured.

E.3.3.2 Model the control structure

Control structures are system models composed of one or more controlled process and its controller, connected through control actions and feedback. The controller applies a control algorithm to determine the control actions to provide, and a process model that represents its approximation to the process being controlled. Problems can occur from any aspect of the control loop. The elements in the control structure can be human, or any technology, including those containing complex electronics/software. Most systems have several overlapping and interacting control loops. A simple example control structure is illustrated in Figure E.1.

Figure E.1 A simple example of a hierarchical control structure
(Source: derived from *STPA Handbook*; N. Leveson, J. Thomas.)

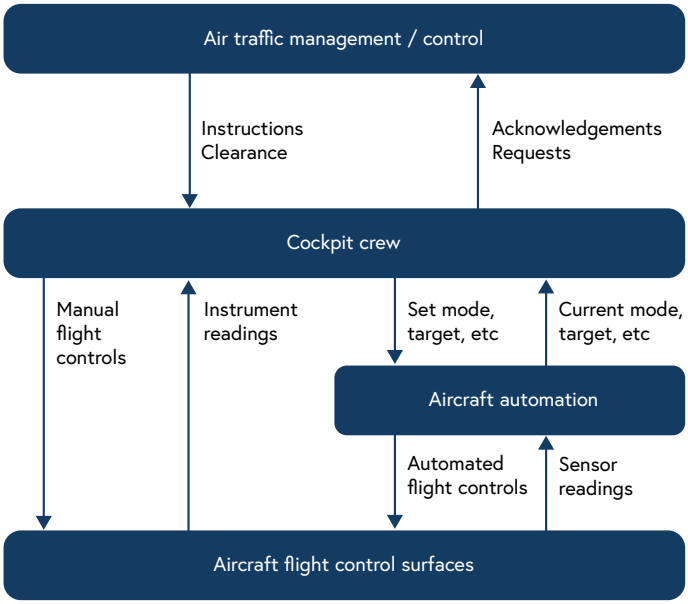

E.3.3.3 Identify Unsafe Control Actions (UCAs)

Unsafe Control Actions (UCAs) are those that in a particular context will lead to a hazard. The STPA process systematically steps through the control actions and identifies potential ways that could be unsafe by consideration of four cases:

1. Not providing the control action leads to a hazard
2. Providing the control action leads to a hazard
3. Providing a potentially safe control action too early, too late, or in the wrong order
4. The control action lasts too long or is stopped too soon

Each UCA is described in five parts:

<source>	the controller that can provide the control action
<type>	not provided; provided; too early/too late; stopped too soon/applied too long
<control action>	the control action or command itself (from the control structure)
<context>	under what conditions the control action is unsafe
<link to hazard(s)>	link to one or more hazards (or sub-hazards)

E.3.3.4 Identify loss scenarios

A loss scenario describes the causal factors that can lead to the UCAs and to hazards. STPA considers why UCAs could occur and why control actions may be improperly executed or not executed at all, leading to a hazard. It does this by addressing:

(a) unsafe controller behaviour;
(b) causes of inadequate feedback/information;
(c) control path; and
(d) other factors related to the controlled process.

E.3.3.5 Output of STPA

The STPA process systematically addresses the potential causes of system hazards and therefore allows requirements to be generated to avoid them or control their effects. Such requirements form part of the design and development process to ensure and assure that they have been adequately addressed.

E.3.4 How it can be adapted

STPA is not inherently a safety analysis method and can be readily applied to security issues alongside safety. The consideration of loss scenarios allows many forms of security threats to be considered, including social engineering of a human controller, attacks on the communications infrastructure forming links in control/feedback loops, and compromise of the logic in the controller or controlled equipment. There are already examples of STPA being applied to safety, security and combined safety/security (see Section E.3.6 below).

E.3.5 How it helps

STPA is designed to facilitate analysis of a wide range of concerns including safety and security, particularly for highly complex and integrated systems with human and complex electronics/software. It enables a common technique to be applied by both cyber security and safety engineers, and allows collaboration between these disciplines in understanding how a cyber attack could cause a safety loss.

E.3.6 Further information

- *STPA Handbook*; N. Leveson, J. Thomas. [https://psas.scripts.mit.edu/home/materials/]
- *Engineering a Safer World: Systems Thinking Applied to Safety*; N. Leveson. [https://mitpress.mit.edu/books/engineering-safer-world]
- *An Integrated Approach to Safety and Security Based on Systems Theory*; W. Young, N. Leveson [http://sunnyday.mit.edu/papers/cacm232.pdf]
- *STPA-SEC for Cyber Security/Mission Assurance*; W. Young [https://psas.scripts.mit.edu/home/wp-content/uploads/2014/03/Young_STAMP_2014_As-delivered.pdf]
- *System-Theoretic Process Analysis for Security (STPA-SEC): Cyber Security and STPA*; W. Young [https://psas.scripts.mit.edu/home/wp-content/uploads/2017/04/STAMP_2017_STPA_SEC_TUTORIAL_as-presented.pdf]
- *STPA-SafeSec: Safety and security analysis for cyber-physical systems*; I. Friedberg, K. McLaughlin, P. Smith, D. Laverty, S. Sezerb [https://www.sciencedirect.com/science/article/pii/S2214212616300850]
- *A Systems Approach to Security: Lessons from the Frontlines Applying STPA-Sec*; W. Young [https://www.acsac.org/2016/program/files/ACSAC2016-SSE-Young.pdf]

E.4 Fault Tree Analysis (FTA)

E.4.1 Introduction

Type: deductive causal analysis, determining the contributory events to an event of interest

The technique works from the event of interest to identify potential causes and conditions that could lead to its occurrence. It is traditionally used for quantitative analysis of failures that lead to a top-level system failure event. In the case of its use to support security considerations, it is its ability to model fault conditions that lead to an undesirable system effect through **qualitative** analysis that is of interest. This has a number of similarities with the use of Attack Trees in security analysis.

Fault Tree Analysis can be used to:

(a) understand the logic leading to the top event/undesired state;
(b) identify critical assets/events;
(c) evaluate sensitivity to certain conditions/faults, aiding determination of acceptable response times to those conditions; and
(d) assist in decomposing high-level fault targets in system design.

An important distinction is made between **fault** and **failure**. A **fault** is taken to be the condition where a system/element/component has an inability to perform as required. It can be due to a hardware **failure** of a component, or from a deficiency such as errors in specification, design, manufacture, maintenance, or its application/operation.

E.4.2 When to apply it

FTA is often applied late in the system's development lifecycle, but it can be applied early on when conceptual architectures and functions are being evaluated. Most benefit can be gained by applying it early in development at a high level of abstraction, and revisiting/further developing the analysis as the design matures.

E.4.3 Basic method

To provide an effective system analysis, a structured method comprising the following elements is conducted:

(a) definition of the purpose and scope of the analysis;
(b) familiarization with the design, functions and operation of the system;
(c) definition of the top event (event of interest);
(d) construction of the fault tree;
(e) analysis of the fault tree logic;
(f) reporting on results of the analysis; and
(g) assessment of implications of the results.

This Section focuses on the construction steps as the basis for adaptation to address safety/security combined objectives.

E.4.3.1 Fault tree construction steps

1 Identification of the event of interest:

In a safety analysis, the top event of interest is usually related to an undesirable event/condition such as a hazard.

2 Identification of **immediate**, **necessary** and **sufficient** causes:

Immediate is meant in a logical rather than temporal sense; for example, turning off a light switch is an immediate (in a temporal sense) cause of a lack of illumination from a filament light bulb; however, in a logical sense, the cause is a lack of electrical current flowing through the filament. Jumping straight to the switch could lead the analyst to overlook a blown bulb as the cause of a lack of illumination. The switch will be identified as a potential cause through a systematic breakdown of the top-level event into intermediate events and basic events.

Necessary is judged in the sense of whether the higher-level event occurs if any of the causal events are not present. Where multiple contributory events are required to cause the higher-level event, they are combined with an AND gate.

Sufficient is judged in the sense of whether the higher-level event will occur with just the identified causal event. Where multiple contributory events could individually cause the higher-level events, they are combined through an OR gate.

Causes are described in concise statements of what the fault is and when it occurs.

3 Classification of intermediate events:

If an immediate cause of the fault is a 'component failure'[56] then the event is classified as a '**state-of-component** fault'; otherwise, it is classified as a '**state-of-system** fault'.

If it is classified as a **state-of-system** fault, then the immediate, necessary and sufficient causes are identified (see step 2). If it is classified as a **state-of-component** fault, then an OR gate is used to combine a standard set of causal events representing *primary*, *secondary* and *command* modes.

Figure E.2 Classic FTA state-of-component structure

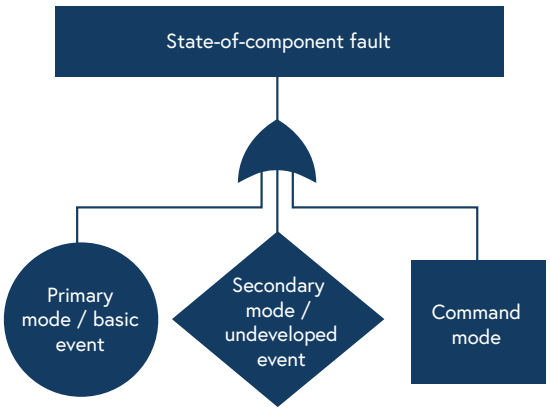

Primary modes relate to failure mechanisms within the component itself and are typically represented by a basic event (one that is not further decomposed); however, the basic event can be decomposed into random physical mechanisms and systematic failure mechanisms combined using an OR gate into the primary fault event. This approach is particularly useful when considering complex or programmable components.

56 It should be noted that in this context, component is interpreted as the lowest level element of consideration in the scope of the analysis. This could be an assembled item, unit or even a system. Failure does not imply any particular mode, state or persistence; it could be due to transient events, specification errors, etc.

Secondary modes relate to failure mechanisms in the environment of the component, such as failure induced by excessive heat caused by an adjacent component failing or by a loss of cooling effect. It is common for secondary modes not to be developed and either represented by an undeveloped event or omitted entirely from the tree structure.

Command modes relate to faults induced by the interfaces to the component.

Figure E.3 Extended FTA state-of-component structure for systematic causes

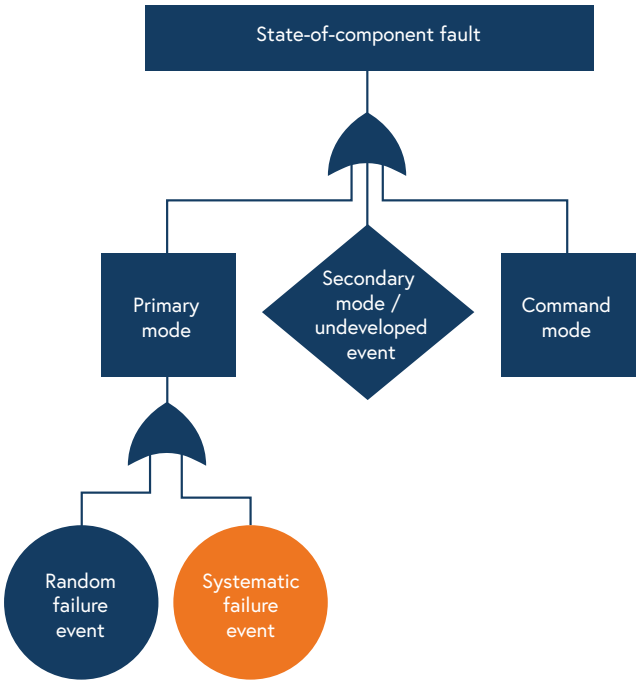

In the earlier example of non-illumination fault, the primary mode would include the 'blown bulb' due to wear out; the secondary mode could be 'blown bulb' due to excessive voltage or excessive vibration/ mechanical shock; and the command mode could be due to lack of voltage across its terminals, that is, it is being commanded not to illuminate.

4 Repeat steps 2 and 3.

Steps 2 and 3 are repeated until all intermediate events are decomposed to basic or undeveloped events. Careful and consistent naming of the basic/undeveloped events allows the analysis algorithms or further common cause/common mode analysis to identify when the same underlying cause contributes to apparently independent branches of the fault tree.

E.4.3.2 Analysis of the fault tree logic

Analysis of the fault tree logic always starts with a qualitative step to produce minimal cut-sets. This step is conducted automatically by FTA tools. The cut-sets are an expression of the combination of basic and undeveloped events that can lead to the event of interest. Boolean logic is used to reduce the total possible combinations represented in the tree logic into the minimal combinations that are enough to cause the top event of the tree. The 'order' of a minimal cut-set identifies how many events are required to cause the top event. A minimal cut-set of order '1' means that a single failure can cause the top event.

By providing information on the failure rates/likelihood of the basic events in the fault tree and the system usage, the likelihood of the intermediate and top events can be calculated to perform a quantitative analysis. This can be misleading if there is a lack of confidence in the figures provided for the basic events and therefore quantitative analysis of a fault tree with systematic events is not recommended.

E.4.4 How it can be adapted

A further extension of the primary mode decomposition can be made to add a cyber security cause. This would represent the case where a security attack is successful in modifying the function of a component. This can start as an undeveloped event and be evaluated during the qualitative analysis of the fault tree logic to determine its significance.

Similarly, the command mode decomposition could consider whether the component can be commanded into a fault state through a malicious action, without modifying the component. Such an action may occur through allowing a command action that is not anticipated in normal use, or provision of inputs that prevent the normal operation of the component, for example, through a denial-of-service attack.

The secondary mode can be used to consider how a security attack can be used to induce stress on a component to cause its premature failure. STUXNET [Ref 9] can be seen as an attack that (in part) used the command mode of one component to induce a secondary mode failure of another by stressing specific components beyond their design limits.

Figure E.4 Extended FTA state-of-component structure for cyber security causes

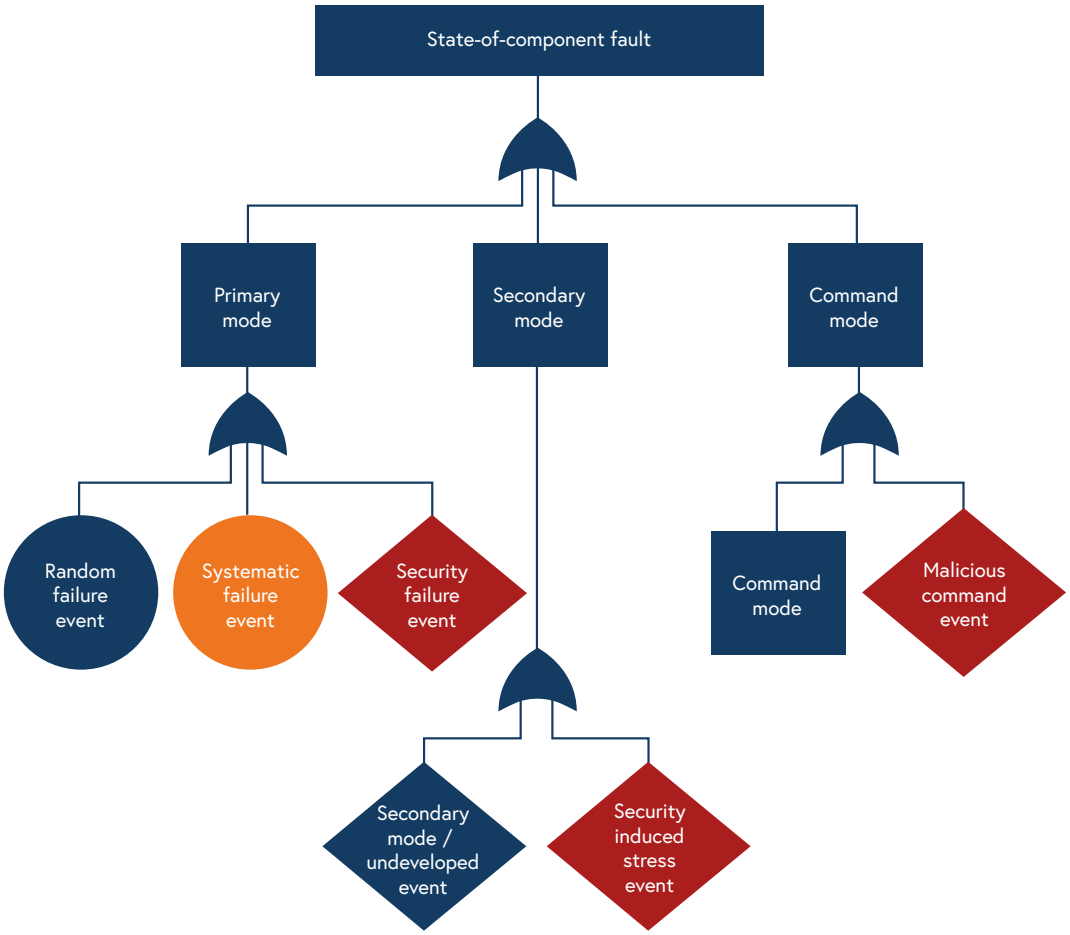

If any such security event appears in a first order minimal cut-set, it identifies that component as a critical item. If any minimal cut-set is made solely of such events, it indicates that a targeted attack may be used to compromise a system to cause a harmful event represented as the top event of the fault tree. The 'security failure event' or set of events can then be investigated further to determine the credibility of the event/event set, and this analysis can be fed back into the safety risk qualitative assessment.

As with the choice of which secondary mode events to develop in a traditional fault tree, judgement will be required to decide which of the security events to develop/investigate to avoid a proliferation of activity out of proportion with the risks. For example, it may be sufficient to add such security events to the development only of components that include programmable elements, as these are the most likely to be vulnerable to security attacks. Alternatively, it may be considered necessary to include critical safety devices that are controlled by programmable elements where the programmable elements control the critical device and therefore could be used to induce excessive stresses that cause premature failure. A joint safety/security review is likely to be the best way to reach this judgement.

E.4.5 How it helps

The approach can be used to aid a shared understanding between security and safety engineers of the relative importance of security events to the safety outcome. It avoids either safety or security disciplines having to understand the methods of the other, whilst being able to clearly define the combination of events that are of mutual interest.

It is emphasized that this approach is intended to inform a qualitative understanding of the security risk to safety. The use of the quantitative methods of FTA are not considered appropriate. This does not prevent these events being included in an FTA that quantitatively analyses system hardware failures, but it should not be portrayed as a quantitative safety/security analysis result.

E.4.6 Further information

- NUREG-0492 *Fault Tree Handbook*: US Nuclear Regulatory Commission [https://www.nrc.gov/docs/ML1007/ML100780465.pdf]
- Fault Tree Analysis (FTA): IEC 61025:2006

E.5 Structured What-If Technique (SWIFT)

E.5.1 Introduction

Type: qualitative risk identification technique: a facilitated brainstorming group activity

SWIFT is simple to use and requires no specialized tools or techniques for the team. Individuals with little hazard analysis training can participate in a full and meaningful way. It can be applied at any time of interest, such as during concept phase, design, operations or maintenance. It can be quicker to apply than HAZOPS or Failure Modes and Effects Analysis (FMEA) and therefore may be seen as a more efficient use of the team's time.

The technique relies heavily on the expertise of and preparation by the facilitator, and the domain experience and intuition of the review team. It is more subjective than some other methods that require a more formal and systematized approach, for example, HAZOPS.

E.5.2 When to apply it

SWIFT can be carried out across a range of system hierarchy levels. It can be applied to processes, architectures and designs. It can therefore be applied when little is known about design implementations and reapplied when details are known, using more focused 'what if' phrases. It is best applied when the level of elements in scope are reduced to a manageable level, either by abstraction or by sub-dividing the scope.

E.5.3 Basic method

The general process is as follows:

1 A suitable prompt list of words or phrases is created: these may be based on a standard set or be created to address the specific case.
2 The context of the item, system, change or situation and the scope of the study are agreed.
3 Discussion is facilitated by creating a question using a 'what-if' phrase and a prompt word or subject to stimulate the study team into exploring potential scenarios, their causes, consequences and impacts.

The 'what-if' phrases may be of the form: "what if…", "what would happen if…", "could someone or something…", "has anyone or anything ever…."

The facilitator asks the participants to raise and discuss known risks and hazards; previous experience and incidents; known and existing controls and safeguards; regulatory requirements and constraints.

4 Risks are summarized and the team considers controls in place.
5 The description of the risk, its causes, consequences and expected controls are confirmed with the team and recorded.
6 The team considers whether the controls are adequate and effective and agree a statement of risk control effectiveness. If this is less than satisfactory, the team further considers risk treatment tasks and potential controls are defined.
7 During this discussion, further 'what-if' questions are posed to identify further risks.
8 The facilitator uses the prompt list to monitor the discussion and to suggest additional issues and scenarios for the team to discuss.
9 Risks may be ranked taking into account the existing controls and their effectiveness and qualitative assessments of residual risk.

E.5.4 How it can be adapted

Over the past few years a number of cyber attacks on safety-involved operating technology have taken place and lessons can be learnt from these in terms of the design and exploitation paths used. As our understanding grows, we can use these attacks to support structured risk assessment. SWIFT is a prospective threat/hazards analysis method that uses structured brainstorming with guide questions, prompts and risks identified. These risks can then be assessed, and mitigations designed and developed.

Table E.2 Examples of 'what if' security-informed supplementary prompts, with example exploits and effects

What ifs?	Example attack	How it was exploited	Effect that was created
Malware was transferred to our system?	German nuclear power plant, malware attack – 2016	W32.Ramnit and Conficker malware on the fuel assembly loading system.	If the infected systems were connected to the internet, they could be activated, changed and data stolen.
Operating system administration services are accessed by malicious actors?	Ukraine energy grid, operational disruption – 2015	Attackers obtained credentials via spearphishing, established remote access via virtual private network, and used administrator services to access the Supervisory Control and Data Acquisition (SCADA) network. (See below for continuation.)	Seven 110 kV and 23 35 kV substations were disconnected for three hours. Later statements indicated that the attack impacted additional portions of the distribution grid and forced operators to switch to manual mode.
Operating technology firmware is overwritten?	See above	Overwriting firmware, followed by accessing breakers and universal power supplies to turn off power and ultimately create an outage.	See above.
Supervisory Control and Data Acquisition systems are accessed by malicious actors?	US Bowman Dam remote access to SCADA systems – 2013	Repeated unauthorized, remote access to SCADA systems controlling the sluice gate, altering water levels and flow rates.	Malicious actor obtained information regarding the status and operation of the dam, including information about the water levels, temperature and status of the sluice gate. This would normally have permitted the remote operation and manipulation of the sluice gate, but the sluice gate had been manually disconnected for maintenance at the time of the intrusion.
Operating safety and security information was made available to malicious actors?	Japan railway operator, data exfiltration – 2015	Spearphishing emails to deliver the Emdivi RAT and then unsuccessfully attempted to take documents regarding railway crime prevention, railway communication systems, safety check procedures, security information and railway safety.	If successful, it would have enabled reconnaissance of transport safety-involved operating technology.
Application-specific malware intrusion detection is not designed in?	Middle Eastern oil and gas petrochemical facility safety system shutdown – 2017	Malware TRITON (also known as TRISIS or HatMan) directly interacting with a Safety Instrumented System (SIS).	Controllers unable to perform tasks like regulating voltage, pressure and temperatures[57].

E.5.5 How it helps

SWIFT is a technique that is familiar to many safety engineers. By adapting the prompts, security experts can be included in the risk identification process. This enables safety and security engineers to work together to achieve 'security-informed safety' and can be used as a precursor to each discipline pursuing their specialist techniques to refine the issues identified, before coming back together to address the safety/security implications.

57 This effect was not realized in the actual attack due to a flaw in the malicious code.

E.5.6 Further information

- ISO 31010:2010, Annex B.9 *Risk management. Risk assessment techniques*
- *A SWIFT response*: IOSH magazine [https://www.ioshmagazine.com/article/swift-response]
- Acquisition Safety & Environmental Management System: ASEMS Toolkit: UK Ministry of Defence (MOD) [https://www.asems.mod.uk/toolkit/swift]
- *Guidance on Hazards Identification*: European Strategic Safety Initiative, ECAST [https://www.easa.europa.eu/sites/default/files/dfu/ECASTSMSWG-GuidanceonHazardIdentification1.pdf]

E.6 Identification of critical digital assets

E.6.1 Introduction

Type: risk treatment – proportionality of controls

There should be one common process for determining criticality of digital assets that can contribute to risks. Some regulatory environments may prescribe a method for identifying critical digital assets, which should be used in place of the method outlined in this Annex. Other regulatory environments set requirements in terms of outcomes and a process such as that described in this Annex may form part of the argument and evidence to support claims of safety and security.

E.6.2 When to apply it

This technique should ideally be applied from concept phase, but can be applied at any point, including for operational/legacy systems.

E.6.2.1 Purpose of identifying critical digital assets

The control systems of a plant/facility can be decomposed into separate digital assets in order to aid safety analysis and security analysis. Some digital assets will be more critical than others in maintaining safe and secure operations and accordingly those digital assets should attract proportionately greater protection – to avoid having to protect every digital asset to match the most stringent requirement. Also, regulations will generally set a boundary for their scope, allowing identification of which functions and digital assets fall within and which fall outside.

Therefore, there may be two reasons for identifying critical digital assets:

1 to assign protection resources to digital assets according to the significance of the function(s) that rely on their correct operation (for example, digital assets that perform, control or support a function); and

2 to identify which functions and consequently which digital assets fall within a regulatory boundary.

E.6.3 Basic method

Despite the goal being to identify the critical digital assets, the starting point must be to understand the functions and interdependencies of the facility/plant. To jump directly to consideration of each digital asset in isolation will risk overlooking interdependencies and common mode failures that could be exploited by a malicious actor, particularly where multiple failures or exploits are orchestrated to maximize their impact.

E.6.3.1 What constitutes a critical function

The facility/plant needs a list of the main functions of the plant that are necessary to meet its business goals and comply with its regulatory restrictions. This includes those functions in the operational technology domain and those functions in the information technology domains that support the business operations, for example, communications, payroll and contact information. Then, for each function, it is necessary to identify the hazardous states that may result from that function, particularly if that function were to operate outside its design envelope. Various process hazard analysis techniques used in safety may be employed here, such as running a HAZOPS.

However, in order to support the identification of wider business risks (for example, loss of delivery of essential services as per the NIS Directive [Ref 11], loss of revenue and cash flow) and the identification of information security risks (for example, a data breach under the General Data Protection Regulation (GDPR) or one that jeopardizes business secrets), the definition and range of hazardous states is likely to need expanding from those that may currently be used for safety analysis.

The output of this analysis should be a table of functions, at a convenient level of granularity, with an assessment of the worst outcome(s) were that function to lead to a hazardous state.

E.6.3.2 Assigning an impact level to that function

A plant/facility should create a graded scale of impact levels, for example, 'High', 'Medium' and 'Low', that calibrates all aspects of risk, including safety-related, business-related, security-related or regulatory sanction, or a risk resulting in a loss to the reputation or profitability of the business.

Each function can be assessed for its criticality if it 'went wrong' and an impact level assigned against that scale of impact. Safety, Security and Emergency Preparedness (SSEP) functions are likely to be assigned to the highest impact level.

Many functions are dependent upon other functions. For example, cooling may require an adequate supply of water and continuity of electric power. If failure to cool for 30 minutes could lead to a hazardous state with a high impact level, the functions to provide water and electric power may inherit that same high impact level unless there is an argument to justify a lower impact level, for example, alternate sources that can provide sufficient cooling using water and/or electric power within 5 minutes may be given a reduced impact level if they are independent.

E.6.3.3 Which digital assets are required for the proper operation of which functions

The control systems of the facility/plant need a decomposition into digital assets using a suitable level of granularity. This should come about from systems engineering and ideally should be designed into the architecture, but could be done retrospectively for a legacy design. A table relating assets to functions should be constructed. That table should identify the following types of relationship:

(a) The digital asset directly **performs** the function. For example, the function is to provide cooling: one of the digital assets that performs the function is a water pump with integrated Programmable Logic Controller (PLC).

(b) The digital asset **controls** the function. For example, the digital asset is a human-machine interface (HMI) that can command the PLC.

(c) The digital asset can **influence** the function. For example, the digital asset is an engineering workstation that is used with many different PLCs and can change set points of the water pump's PLC. Alternatively, the digital asset is another workstation on the same network that the PLC trusts.

(d) The digital asset *supports* the function. For example, the digital asset is an automated switch panel that provides electric power to the pump and many other digital assets. Note that the panel also **performs** part of the plant's electric power function, which *supports* the cooling function. Alternatively, the digital asset provides network timing services on which the water pump's PLC relies.

E.6.3.4 Assigning digital assets to impact levels

Each digital asset should then be assigned to an impact level according to these criteria:

(a) For each digital asset, it inherits the highest impact level of the function(s) that it performs.

(b) For each digital asset, we need to identify any trusted relationship with another digital asset, such that the first **supports**, can directly **control** or *influence* the operation of the second. In this case, in the absence of an argument that there is a practical means of intervention to stop the propagation of an error condition or of malicious action, the digital assets must collectively inherit the highest impact level of the function(s) that they perform or support. Clearly, there is a strong motive here to choose measures that limit the propagation of errors or malicious action, because otherwise the whole plant/facility will need to be protected to the most stringent level required anywhere in that plant.

(c) Digital assets that are identified as providing an independent layer of protection in safety analysis, for example, through a Layers of Protection Analysis (LOPA), may warrant a higher impact level, determined by the impact of the hazardous state that the asset protects against, modified by the safety claim of its strength, for example, its reduction of the probability of failure on demand. Company policy needs to create the appropriate tables for this calculation[58].

E.6.3.5 Determining which assets are deemed critical

Digital assets that fall within the regulatory boundary may be labelled as critical digital assets. The determination will depend upon the nature of the regulation.

Outside consideration of regulation, the label 'critical digital asset' should be defined in a way that is useful to the company: the expression might not be used. Digital assets that are assigned a sufficiently high impact level should be deemed critical. In a two-level scale, the digital assets that perform the higher impact level functions will be labelled 'critical digital assets'. In a scale of three impact levels, it may be just the highest impact level that applies. In a scale of five levels, it may be the top two or three that apply.

E.6.3.6 Protecting those digital assets in a proportionate way

The plant/facility should define in its policy the requirements for protective measures for each impact level, to achieve a graded approach with the highest impact levels benefiting from the most stringent and consequently most expensive measures. For example, physical security measures may include an out-of-hours alarm for the lowest impact level, but patrolling security guards and 24-hour CCTV for the highest. Cyber security measures may specify different levels of system hardening and network hardening, for example, zoning. Organizational and procedural measures may include the nature of personnel background checks and supply chain requirements imposed for the different impact levels.

There are no standards specifying what protection is sufficient, because this depends on assumptions about the nature of the threat, vulnerabilities, impact and the risk appetite of the organization and its regulator.

E.6.3.7 Identifying additional requirements for the protection of digital assets

Analysis of failure modes of the digital asset, either as a result of a random component failure or as a result of malicious action orchestrating multiple simultaneous failures, will identify one of the following conditions for the function with which it is associated:

1 The function and associated assets continue to operate as designed, i.e., The function is resilient.

2 The function and associated assets operate in an observable manner in a region of extended operation; for example, the pump is running tolerably too hot as long as timely action is taken.

3 The function fails in a predictable and observable manner, i.e., an anticipated failure mode.

4 The function fails in an observable but unpredictable manner.

5 The function fails in an unobservable manner.

58 At the time of publication, no accepted standards for this calculation are known.

The conditions are in order of increasing risk to the plant/facility and may require additional measures to be applied to the function and its associated assets to mitigate the impact. These measures will be in addition to those specified by policy resulting from the assignment to an impact level. For example, if the asset is the PLC that controls the water pump, described earlier, analysis of the control logic may reveal that the failure of a single sensor results in condition 2 or 3, but the malicious spoofing of two sensors may lead to condition 4 or 5. The security measure may be to protect the sensors or modify the control logic.

It is possible that these additional requirements will identify particular digital assets as warranting assignment to a higher impact level or additional measures specifically for that digital asset. It is important that the standard requirements for digital asset protection at a particular security level represent the norm rather than the extreme cases.

E.6.3.8 Validating that the right digital assets have been identified

Once the digital assets have been identified and assigned to impact levels, the policy has been applied and protective measures have been allocated to each asset, the defensive design should be tested against a representative range of malicious attack scenarios. This analysis should be performed in the manner of a 'red team' review, to identify weak points in the design and/or implementation of the defences. This analysis may identify a particular digital asset as being highly attractive to an attacker because it is a common enabler for many different attack scenarios: for example, a camera that monitors a choke-point. That digital asset may warrant reassignment to a higher impact level.

E.6.4 How it can be adapted

The above approach is adapted from common risk management methodologies by devising a common risk criteria framework.

E.6.5 How it helps

The use of common risk criteria aids shared understanding of the risk to the operation from all domains and thus helps to achieve a proportionate response.

E.6.6 Further information

- *None*

Annex F

Bibliography

This Annex lists standards/articles, etc. that are referenced by this Code. It also identifies additional material that has been consulted in the authoring of this Code or that may provide a source of further information for those addressing the recommendations of this Code.

Additional references specific to the techniques and measures discussed in Annex E are listed at the end of each section within that Annex.

Website references were last accessed and correct at August 2020.

F.1 Referenced documents

[Ref 1]	*Security-Informed Safety: If It's Not Secure, It's Not Safe*: R. Bloomfield (Centre for Software Reliability), K. Netkachova (City University London), R. Stroud (Adelard LLP) [https://www.adelard.com/assets/files/docs/Bloomfield_serene_2013.pdf]
[Ref 2]	IEC 61508 Parts 1-7, (2010) *Functional safety of electrical/electronic/programmable electronic safety-related systems*
[Ref 3]	RTCA DO-356A/EUROCAE ED-203A *Airworthiness Security Methods and Considerations*
[Ref 4]	HC 787; 18/04/2018: *Cyber-attack on the NHS* [https://publications.parliament.uk/pa/cm201719/cmselect/cmpubacc/787/787.pdf]
[Ref 5]	*Computer Bugs In Hospitals: A New Killer*: Prof. M. Thomas, Prof. H. Thimbleby; 06/02/2018, Gresham College [https://s3-eu-west-1.amazonaws.com/content.gresham.ac.uk/data/binary/2642/2018-02-06_MartynThomasHaroldThimbleby_ComputerBugs.pdf] [https://www.gresham.ac.uk/lectures-and-events/computer-bugs-in-hospitals-a-new-killer]
[Ref 6]	*Remote Exploitation of an Unaltered Passenger Vehicle*: Dr C. Miller, C. Valasek; 10/08/2015 [http://illmatics.com/Remote Car Hacking.pdf]
[Ref 7]	*Why Car Hacking Is Nearly Impossible*: D. Pogue; 28/10/2016, Scientific American [https://www.scientificamerican.com/article/why-car-hacking-is-nearly-impossible/]
[Ref 8]	*MAR-17-352-01: Hatman-Safety System Targeted Malware (Update B); 27/02/2019*; 18/12/2017, US National Cybersecurity and Communications Integration Center, US Dept of Homeland Security [https://us-cert.cisa.gov/ics/MAR-17-352-01-HatMan-Safety-System-Targeted-Malware-Update-B]
[Ref 9]	*W32.Stuxnet Dossier*: N. Falliere, L. O. Murchu, E. Chien; Version 1.4 (February 2011) [https://www.wired.com/images_blogs/threatlevel/2011/02/Symantec-Stuxnet-Update-Feb-2011.pdf]
[Ref 10]	*Framework for Improving Critical Infrastructure Cybersecurity*; Version 1.1 16/04/2018, National Institute of Standards and Technology [https://www.nist.gov/cyberframework] [https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf]
[Ref 11]	EU: Directive 2016/1148: concerning measures for a high common level of security of network and information systems across the Union (The Directive on security of network and information systems – NIS Directive); 06/07/2016 UK: Statutory Instruments 2018 No. 506: The Network and Information Systems Regulations 2018 [https://www.legislation.gov.uk/uksi/2018/506/made] [https://www.gov.uk/government/collections/nis-directive-and-nis-regulations-2018] [https://ico.org.uk/for-organisations/the-guide-to-nis/]
[Ref 12]	ISO/IEC 27000 family: Information security management systems
[Ref 13]	Cyber Essentials [https://www.ncsc.gov.uk/cyberessentials]
[Ref 14]	IEC 62264-1:2013 *Enterprise-control system integration – Part 1: Models and terminology*
[Ref 15]	ISO/IEC/IEEE 15288:2015 *Systems and software engineering – System life cycle processes*

Annex F – Bibliography

[Ref 16]	SafSec was a UK Ministry of Defence (MOD) project developing a common methodology for security accreditation and safety assurance. The reports from this project are no longer available online, but papers exploring the approach are still available: *SafSec: Commonalities Between Safety and Security Assurance*. Thirteenth Safety Critical Systems Symposium, Southampton: S. Lautieri, D. Cooper, D. Jackson; 2005 [https://scsc.uk/scsc-7] ISBN: 1-85233-952-7 [https://link.springer.com/chapter/10.1007/1-84628-130-X_5] *SafSec: Combining Security and Safety Principles in Practice*. Second Institution of Engineering and Technology International Conference on System Safety: T. Cockram, S. Lautieri; 2007; pp159-164. ISBN: 978-0-86341-863-1.
[Ref 17]	ISO 31000:2018 *Risk management – Guidelines*
[Ref 18]	*STPA Handbook*; March 2018 [https://psas.scripts.mit.edu/home/get_file.php?name=STPA_handbook.pdf]
[Ref 19]	*Out of control: Why control systems go wrong and how to prevent failure*; 2003 [https://www.hse.gov.uk/pubns/priced/hsg238.pdf]
[Ref 20]	SCSC-159: *Assurance Case Guidance 2020*: SCSC Assurance Case Working Group [https://scsc.uk/scsc-159]
[Ref 21]	*The role of hierarchical knowledge representation in decision making and system management*: J. Rasmussen; 1970; IEEE Transactions on Systems, Man & Cybernetics; SMC-15. 234-243. 10.1109/TSMC.1985.6313353 [https://ieeexplore.ieee.org/document/6313353]
[Ref 22]	*Systems of Systems Primer*: INCOSE [https://www.incose.org/products-and-publications/sos-primer]
[Ref 23]	*Systems Security Engineering – Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems*: NIST Special Publication 800-160, Volume 1 [https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160v1.pdf]
[Ref 24]	*Developing Cyber Resilient Systems: A Systems Security Engineering Approach*: NIST Special Publication 800-160, Volume 2 [https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160v2.pdf]
[Ref 25]	HSE OG-0086; Edition 2: *Cyber Security for Industrial Automation and Control Systems*: [https://www.hse.gov.uk/foi/internalops/og/og-0086.pdf]
[Ref 26]	*Nuclear Safety Technical Assessment Guide: Computer Based Safety Systems*: ONR NS-TAST-GD-046, Revision 5 [http://www.onr.org.uk/operational/tech_asst_guides/ns-tast-gd-046.pdf]
[Ref 27]	*Security Assessment Principles (SyAPs)*: ONR [http://www.onr.org.uk/syaps/index.htm]
[Ref 28]	*The Precautionary Principle*: World Commission on the Ethics of Scientific Knowledge and Technology (COMEST), United Nations Educational, Scientific and Cultural Organization (UNESCO); 2005 [https://unesdoc.unesco.org/ark:/48223/pf0000139578/PDF/139578eng.pdf.multi] [https://unesdoc.unesco.org/ark:/48223/pf0000139578]
[Ref 29]	Communication from the European Commission on the Precautionary Principle /* COM/2000/0001 final */ [https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=LEGISSUM:l32042]
[Ref 30]	ISO 28000:2007 *Specification for security management systems for the supply chain*
[Ref 31]	*Independent Co-Assurance using the Safety-Security Assurance Framework (SSAF): A Bayesian Belief Network Implementation for IEC 61508 and Common Criteria*: N. Johnson, Y. Gheraibia, T. Kelly The paper is available from [https://scsc.uk/rp154.13:1/] as part of the symposium proceedings [https://scsc.uk/scsc-154]
[Ref 32]	ISO 55000:2014 *Asset management – Overview, principles and terminology*
[Ref 33]	*Cybersecurity of medical devices – Addressing patient safety and the security of patient health information*: R. Piggin [https://www.bsigroup.com/LocalFiles/EN-AU/ISO 13485 Medical Devices/Whitepapers/White_Paper___Cybersecurity_of_medical_devices.pdf]

Annex F – Bibliography

F.2 Additional reading

F.2.1 Standards

- BS 10754-1:2018 *Information technology. Systems trustworthiness. Governance and management specification*
- BS 31111:2018 *Cyber risk and resilience. Guidance for the governing body and executive management*
- PAS 11281:2018 *Connected automotive ecosystems. Impact of security on safety. Code of practice*
- ISA-TR84.00.09-2017 *Cybersecurity Related to the Functional Safety Lifecycle*
- ED-201 *Aeronautical Information System Security (AISS) Framework Guidance*
- ED-203A/RTCA DO-356A *Airworthiness Security Methods and Considerations*
- ED-204/RTCA DO-355 *Information Security Guidance for Continuing Airworthiness*
- ED-202A/RTCA DO-326A *Airworthiness Security Process Specification*
- IEC 62859:2016+A1:2019 *Nuclear power plants – Instrumentation and control systems – Requirements for coordinating safety and cybersecurity*

Where to obtain standards:

Prefix	Expansion	Standards organization	Website
BS	British Standard	BSI	[https://shop.bsigroup.com/]
DO		RTCA	[https://my.rtca.org/nc__store]
ED	EUROCAE Document	EUROCAE	[https://eshop.eurocae.net/]
IEC	International Electrotechnical Commission	IEC	[https://www.iec.ch/]
ISA	International Society of Automation	ISA	[https://www.isa.org/store]
ISO	International Organization for Standardization	ISO	[https://www.iso.org/]
NUREG	Nuclear Regulation	US NRC	[https://www.nrc.gov/]
PAS	Publicly Available Specification	BSI	[https://shop.bsigroup.com/]

F.2.2 Papers/publications

- *Dependability Terminology: Basic Concepts and Taxonomy of Dependable and Secure Computing*: A. Avižienis, J.C. Laprie, B. Randell, C. Landwehr; IEEE Transactions on Dependable and Secure Computing; 2004. 1(1): p. 11-33.[https://www.nasa.gov/pdf/636745main_day_3-algirdas_avizienis.pdf] [https://ieeexplore.ieee.org/document/1335465]
- *Cyber primer*: UK Ministry of Defence (MOD) [https://www.gov.uk/government/publications/cyber-primer]
- *Availability of Open Source Tool-Sets for CNI-ICS*: RITICS; 23/3/2018 [http://ritics.org/wp-content/uploads/2018/07/Open-Source-Tools-for-ICS-final.pdf]
- DNVGL-RP-G108: *Cyber security in the oil and gas industry based on IEC 62443* [https://rules.dnvgl.com/docs/pdf/DNVGL/RP/2017-09/DNVGL-RP-G108.pdf]
- DNVGL-RP-0496: *Cyber security resilience management: Managing cyber security risks in maritime and offshore industries for ships and mobile offshore units in operation* [https://www.dnvgl.com/maritime/dnvgl-rp-0496-recommended-practice-cyber-security-download.html]
- DNVGL-CP-0231: *Cyber security capabilities of control system components* [https://rules.dnvgl.com/docs/pdf/DNVGL/CP/2018-01/DNVGL-CP-0231.pdf]
- DNVGL-RP-D201: *Integrated software dependent systems* [https://rules.dnvgl.com/docs/pdf/DNVGL/RP/2017-07/DNVGL-RP-D201.pdf]
- *Hunting and Responding to Industrial Intrusions*: Dragos [https://dragos.com/wp-content/uploads/2017-Review-Hunting-and-Responding-to-Industrial-Intrusions.pdf]

Annex F – Bibliography

- *Industrial Control System Threats*: Dragos [https://dragos.com/wp-content/uploads/2017-Review-Industrial-Control-System-Threats.pdf]
- *Industrial Controls System Vulnerabilities: Year in Review 2018*: Dragos [https://dragos.com/wp-content/uploads/yir-ics-vulnerabilities-2018.pdf]
- *NIS Directive and the Security of Critical Services*: Dr R. Piggin [https://www.researchgate.net/publication/323198201_NIS_Directive_and_the_Security_of_Critical_Services]
- *Cybersecurity and Cyber-Resilient Supply Chains*: H. Boyes [https://timreview.ca/sites/default/files/article_PDF/Boyes_TIMReview_April2015.pdf]
- *Code of Practice: Cyber Security for Ships*: IET Standards [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/642598/cyber-security-code-of-practice-for-ships.pdf]
- *Code of Practice for Cyber Security in the Built Environment*: IET Standards [https://shop.theiet.org/code-of-practice-for-cyber-security-in-the-built-environment]
- *Over 20 years of research into cybersecurity and safety engineering: a short bibliography*: S. Paul and L. Rioux [https://www.witpress.com/Secure/elibrary/papers/SAFE15/SAFE15029FU1.pdf]
- *National Cyber Security Strategy 2016 to 2021*: Gov.UK [https://www.gov.uk/government/publications/national-cyber-security-strategy-2016-to-2021]
- *Process safety and cyber security convergence: Lessons identified, but not learnt?*: Dr R. Piggin [https://digital-library.theiet.org/content/conferences/10.1049/cp.2013.1699]
- *Consequence-driven cyber-informed engineering (CCE)*: OSTI.GOV [https://www.osti.gov/biblio/1341416]
- *An Assurance Framework for Independent Co-assurance of Safety and Security*: N. Johnson and T. Kelly [https://www.researchgate.net/publication/327160234_An_Assurance_Framework_for_Independent_Co-assurance_of_Safety_and_Security]
- GAO-19-128: *Weapon Systems Cybersecurity: DOD Just Beginning to Grapple with Scale of Vulnerabilities*: United States Government Accountability Office; Oct 2018 [https://www.gao.gov/assets/700/694913.pdf]
- *Cybersecurity of medical devices: Addressing patient safety and the security of patient health information*: Dr R. Piggin [https://www.bsigroup.com/en-GB/medical-devices/resources/whitepapers/] [https://www.bsigroup.com/LocalFiles/EN-AU/ISO 13485 Medical Devices/Whitepapers/White_Paper___Cybersecurity_of_medical_devices.pdf]
- *Securing critical services* (introducing the Network and Information Systems (NIS) Directive): Dr R. Piggin [https://www.bcs.org/content-hub/securing-critical-services/]
- *Governance, risk and compliance: impediments and opportunities for managing operational technology risk in industrial cyber security and safety*: Dr R. Piggin [https://digital-library.theiet.org/content/conferences/10.1049/cp.2014.0982]

F.2.3 Web resources

- Research Institute in Trustworthy Inter-connected Cyber-physical Systems (RITICS) [https://ritics.org/]
- Partnership for Systems Approaches to Safety and Security (PSASS) [https://psas.scripts.mit.edu/home/]
- Cyber Security Framework: US Government [https://www.nist.gov/cyberframework]
- The National Cyber Security Centre: UK Government [https://www.ncsc.gov.uk/]
  - *Cyber Assessment Framework (CAF) guidance* [https://www.ncsc.gov.uk/collection/caf]
  - *Introducing component-driven and system-driven risk assessments* [https://www.ncsc.gov.uk/collection/risk-management-collection/component-system-driven-approaches/introducing-component-driven-and-system-driven-risk-assessments]
  - *Secure design principles* [https://www.ncsc.gov.uk/collection/cyber-security-design-principles]
- *Security and Safety Modelling*: EU [http://sesamo-project.eu/]
- *Spire: Intrusion-Tolerant SCADA for the Power Grid*: Johns Hopkins University [http://www.dsn.jhu.edu/spire/]
- The Aqua Book (guidance on producing quality analysis for government): UK Government [https://www.gov.uk/government/collections/aqua-book-resources]

Index

Index

Index

Code of Practice

Cyber Security and Safety

This Code of Practice is written for engineers and engineering management to support their understanding of the issues involved in ensuring that the safety responsibilities of an organisation are addressed, in the presence of a threat of cyber attack. *"If it's not secure, you can't be confident it's safe"*.

The implementation of effective cyber security will, in general, require modification of safety-related systems and current procedures throughout their lifecycle. A close interaction between respective engineers is therefore vital. However, teams responsible for safety and cyber security are often in different parts of an organisation. In many organisations, the governance of the combined risk only comes together at a point of such seniority that the technical competence and capacity for detail may be inadequate to ensure that the teams work together effectively. Consequently, the combined risk to the enterprise is not always fully comprehended. Any divergence or conflict between safety and security goals requires the business to make a conscious decision on how to proceed.

The aim of this Code is to help safety-related system practitioners manage cyber security vulnerabilities that lead to hazards. It does this by setting out principles, based on a systems engineering approach, which, when applied, will improve the interaction between the disciplines of functional safety and cyber security, which have historically been addressed as distinct activities.

ISBN 978-1-83953-318-1

9 781839 533181

The Institution of Engineering and Technology
Michael Faraday House
Six Hills Way
Stevenage
Herts
SG1 2AY

theiet.org/standards