ARTICLE:

Duty of care and engineering functional-safety standards

By **Peter Bernard Ladkin**

The operation of some engineered systems may result in harm; a railway level-crossing control, say, which allows cars and trains through, but, one hopes, not at the same time; or a high-pressure pipe with high-temperature contents which could rupture and spray a nearby person. Such systems are designated "safety-related" or "safety-critical" in engineering, and there are standards nominally expressing best practice in their development and operation. One such standard is IEC 61508:2010 (IEC 61508).

The design and operation of safety-related systems is also governed by law. In English law and in common-law-based legal systems, there is a principle of "duty of care" (deriving from *Donoghue v Stevenson* [1932] UKHL 100). Further, if personal harm is caused during the operations of a company, then various employees of the company may be held to be criminally responsible. For example, the Corporate Manslaughter and Corporate Homicide Act 2007 established a criminal offence of "corporate manslaughter", holding a company's directors personally responsible for a death caused as a result of company operations. Safety through corporate operations is also regulated through the Health and Safety at Work etc. Act 1974, which established supervision of such operations in England through the Health and Safety Commission and the Health and Safety Executive (HSE).

This paper considers the interactions between engineering standards and legal principles, but its author is not a lawyer and the paper does not consider the position in English law in any detail. Its purpose is to try to foster a deeper awareness of the issues.

Generally, criminal responsibility for harm may be established if "gross negligence" of an agent is proved "beyond reasonable doubt". Sometimes, the law establishes a "strict liability", in which state of mind[1] elements such as negligence play no role, and all that must be shown "beyond reasonable doubt" is that the act which caused harm was committed by the accused.

A company and its employees have a "duty of care" generally to other persons that they shall not be harmed. If company operations may result in harm, as it might if the company designs, builds or operates safety-critical systems, that duty of care applies.

In England, the Health and Safety Executive brings enforcement actions, including criminal prosecutions, against companies and their officers judged to have failed in their duty of care, as a result of accidents and other operations where someone is harmed (Thomas 2019). In 1997, the International Electrotechnical Commission published the first edition of IEC 61508, the international standard for functional safety of systems involving electrical, electronic or programmable electronic (E/E/PE) subsystems and functional units (IEC 61508). Since the late 1990's, HSE has used conformance/nonconformance with IEC 61508 as a guideline for bringing enforcement actions in the case of harm caused through operation of such systems (Bowell 1997/1999).[2] It is commonplace in electrotechnology that engineered systems are increasingly employing E/E/PE subsystems of some sort, to the point of almost-ubiquity. Use of IEC 61508 and its "industry-specific" derivatives is concomitantly broadening. Indeed, IEC 61508 is classified as a "basic safety standard" by the IEC (Bell 2019).

One of the established principles in safety-related system operations is that a hazard and risk analysis be performed (ISO/IEC Guide 51). Generally speaking, an attempt must be made to identify all the ways – states of the system and its environment; events – in which harm may possibly be caused through system operations. Hazards are, broadly speaking, precursors of harmful events. Identifying them is Hazard

[1] Also referred to as 'mens rea'.

[2] Various HSE officers have said this, to the author and publicly to system-safety specialists, on a number of occasions since the publication of the first edition of IEC 61508. The author recalls Mark Bowell making his statement to a publicly-archived mailing list, the High-Integrity Systems List, hosted then at the University of York. However, he has not been able to trace the archive at time of writing, hence the limited reference.

Identification. Then, the possible harmful circumstances which might follow the hazard, called the hazard severity, are identified. After that, an attempt must be made to assess the likelihood that one of those pathways to harm will actually occur: the likelihood that the hazard will arise; the likelihood that a harmful event of the given severity will arise, and so on. This is risk assessment, risk being defined as a combination of likelihood with severity of harm in electrotechnical standards.[3] It is generally accepted that you cannot usually reduce a risk to zero, but you must reduce it to a level deemed "acceptable". That "acceptable level" is taken to be an increase in chance of death of one in a million per year (HSE 2001, paragraph 130).

When a company engaged in operations that may result in harm has not performed adequate hazard and risk assessment, then it and (some of) its officers may be prosecuted if harm ensues. "Adequate" here usually means state-of-the-practice; over the last century, engineering science has developed methods of performing such hazard and risk assessments. Whatever methods are used, a best effort must be made to identify all the hazards and assess the risk.

Harm may be caused if someone breaks in to the control room of an industrial plant, takes control and causes the plant to operate in a harmful manner. Such "hijacking" must obviously be hindered; physical security in such plants is ubiquitous. More recently, as many such plants have become digitised, and their operation and control of that operation may be effected remotely, through digital communications channels connecting the plant to networks such as the data communications systems run by telephone companies, such hijacking may happen through those channels. Malware programs, which cause possibly harmful and certainly unwanted operations, may be inserted into digital control systems via outside communication channels. Or malware can be left lying around on portable storage, such as a USB stick, in the hope that someone will unthinkingly plug it in to a system USB port at some point. This is becoming almost daily news.[4] The duty of care is no longer discharged merely through erecting fences and having guards who limit physical access and check personal credentials. There needs to be some cybersecurity, as it is now called, some guard on digital communications and against the subversion of plant-critical digital systems, also.

As noted above, HSE has used the E/E/PE functional safety standard IEC 61508 for some two decades as the touchstone for deciding to bring enforcement action in cases in which harm has been caused by operation of a safety-related system involving E/E/PE kit. For example, subclause 7.4.2.3 of IEC 61508:1997 (the first edition) says:

> The hazards and hazardous events of the EUC and the EUC control system shall be determined under all reasonably foreseeable circumstances (including fault conditions and reasonably foreseeable misuse). …….

This clause says you must identify the hazards. Some technical terminology is involved:

> "hazardous event" is an event which may result in harm
>
> "shall" means "must": it is required
>
> "reasonably foreseeable misuse" means use not intended by the operator, but which can arise through "common human behaviour".

There is, though, a problem with "reasonably foreseeable misuse" when you try to fit it to "cyberinsecurity". Vulnerabilities in digital equipment can be quite sophisticated, and require person-months or -years of hard work, and maybe some luck, to discover. The question then arises whether the exploitation of such a vulnerability "arise through common human behaviour"? It can actually be pretty sophisticated behaviour only available to a few criminally-talented people through hard work.

This weakness of phraseology was recognised when the second edition of the standard was prepared. IEC 61508-1:2010 subclause 7.4.2.3 says:

> The hazards, hazardous events and hazardous situations of the EUC and the EUC control system shall be determined under all reasonably foreseeable circumstances (including fault conditions, reasonably foreseeable misuse and malevolent or

[3] There are other notions of "risk" in related engineering disciplines. This is the conception in the International Electrotechnical Vocabulary, IEC 60050.

[4] For example, a renowned southern-German engineering-sensor fabricator at which the author has a number of colleagues was infiltrated in mid-October 2019, at the time this document was being prepared, by malware that had reputedly taken down much of their company IT, including their WWW site, which they replaced with simple HTML code explaining what had happened (Pilz 2019). The E/E/PE subsystems within their products were not affected.

> unauthorised action). …. If the hazard analysis identifies that malevolent or unauthorised action, constituting a security threat, as being reasonably foreseeable, then a security threats analysis should be carried out.

Specifically, "malevolent or unauthorised action" is now included as well as "reasonably foreseeable misuse". So that takes care of that, doesn't it?

Well, not really. Let us parse the phrase using a form of logical controlled-English. Let us call "hazards, hazardous events and hazardous situations" simply "hazards". Let us translate "determined" as "identified". The first statement of the subclause becomes:

(ALL) hazards SHALL be identified

AND

hazards ARE *((reasonably foreseeable fault conditions)*

OR

(reasonably foreseeable misuse)

OR

(reasonably foreseeable (malevolent or unauthorised action))

OR

(other reasonably foreseeable hazards))

It is surely reasonable to require all hazards be identified, but notice that the "*ALL*" is qualified through the second conjunct. The first conjunct does not explicitly say "*ALL reasonably foreseeable hazards SHALL be identified*" but the clause implies this through the use of the qualifier in each listed hazard type. It follows, amongst other things, that the hazard identification process is to take into account not all cyberinsecurities, but those whose exploitation is "reasonably foreseeable". And here there is an issue.

What is "reasonably foreseeable"? Was a Stuxnet-type malware infiltration "reasonably foreseeable" in 2009? Many would say not. Is it "reasonably foreseeable" now that it has happened? Many would say yes, because it exists, and maybe it has been reverse-engineered and distributed more broadly, perhaps in modified form, as malware often is. Suppose you have an unassuming little collection of centrifuges performing some routine but safety-related tasks for your company. Is it "reasonably foreseeable" that it would happen to you? The reasons above may argue for yes. But the malware is popularly held to have been developed by sophisticated government agencies or government-directed organisations of a nation-state, for which your unassuming little collection may be supposed to hold no interest whatever, and that may argue for no. What about the Triton malware (PAS Global, 2019)? It has been used more than once, suggesting that it is "reasonably foreseeable" that it will be used again. A state is also popularly supposed to be involved. But there is a common view in the industry that "there is nothing you can do against that" sort of attack (Anon 2019-1).[5] What kind of duty of care is implied in a situation in which such an attack is reasonably foreseeable but "there is nothing you can do"?

If a "malevolent or unauthorised action" is "reasonably foreseeable", as we have argued above all such identified threats are to be, then subclause 7.4.2.3 says further that you are to perform a "security threats analysis", whatever that is. There seems to be as yet little engineering agreement on exactly what such an analysis is, or indeed what it should be called (Ladkin 2020). And the subclause does not require that the analysis be performed ("*should*" means recommended; it is weaker than "*shall*", which means required, according to IEC conventions on vocabulary). If you perform this analysis and it "identifies security threats", then subclause 7.5.5.2 (below) says to go further.

The reader may now have formed an opinion that the writing in some international engineering standards does not necessarily bear well the intense scrutiny which legal writing and argument typically undergo. I suggest this impression is apt. But recall that IEC 61508 is a touchstone for legal action in connection with a duty of care, and even gross negligence.

It makes sense to interpret the general intent of the subclause in light of a duty of care: that the duty of care requires that the company make a best-effort

[5] As a specialist in software-based system dependability, the author begs to differ substantially with this assessment. That said, the author acknowledges that the management of industrial process plant is a difficult logistical enterprise and to manage such systems perfectly may well be beyond the current capacity of any organisation other than government agencies with potentially unlimited resources.

attempt to identify hazards, including those which can arise through cyberinsecurity, and then do something about it (diminish the vulnerability; inhibit exploitation of it; and so on).

The "doing something about it" is also problematic as handled by subclause 7.5.5.2:

> If security threats have been identified, then a vulnerability analysis should be undertaken in order to specify security requirements.

Consider that you have seen a security threat, maybe you have or you have not performed some analysis of that threat. This subclause recommends you perform a "vulnerability analysis" (whatever that is; the same problem arises as with "security threats analysis") and formulate "cybersecurity requirements". So if you have identified "reasonably foreseeable" "malevolent or unauthorised action", you are to perform a "security threats analysis", and if you identify threats, a further "vulnerability analysis" and formulation of "security requirements". That is, unfortunately, all. The standard does not explicitly say you must do anything to annul or mitigate the threat.

This is unsatisfactory. Your duty of care says you must identify hazards arising through cybersecurities (to state-of-the-practice), and you must do something about them. IEC 61508-1:2010 omits to say either that the "cybersecurity requirements" must address the security threats you identified, or that those requirements must be fulfilled when you deploy your system.

Both should happen. Fixing this omission in the standard would be straightforward. Change those "should"s to "shall"s; replace the word-soup concerning "analysis" with some well-understood term for well-understood analysis; say that the cybersecurity requirements must address the vulnerabilities you found, and require those cybersecurity requirements to be fulfilled in your system. Those conditions can be achieved in about that number of words and there is a proposal (let me call it Proposal 1) to do exactly that for the next edition of IEC 61508 (Ladkin 2018).

On the other hand, the IEC 61508 Maintenance Teams have been presented also with a proposal (Proposal 2) to remove cybersecurity considerations completely from hazard and risk analysis. That you would determine hazards and hazardous events under reasonably foreseeable circumstances, including fault conditions and reasonably foreseeable misuse, *but explicitly not those hazards or hazardous events which result from malevolent or unauthorised action*. In these circumstances, you would mitigate those hazards and hazardous events you identified, as required by the standard, but not necessarily those resulting from malevolent or unauthorised action, because you have excluded them from consideration.

Such a proposal has some support amongst practicing engineers. To try to explain why, I suggest we distinguish between what reasons are given overtly for such a suggestion, and what reasons there may be in the world. Overt reasons are that "cybersecurity" and "safety" are two different qualities of systems for which it makes sense to have two different (sets of) documents describing how to achieve them. But even if one considers that to be so, it is also well-understood that cybersecurity considerations and safety considerations interact, and it is being proposed to remove considerations addressing that interaction from a guidance document. A justification for such a move seems to this author to be lacking. More proximate organisational reasons may be that, in engineering industry, safety departments predominantly consist of specialist engineers, and security departments have traditionally been concerned largely with physical security – fences, locks and locked gates, flashlights, radios, guard-dogs and such like – and in many organisations these two departments remain ill-equipped to work closely with each other. More narrowly, it has also been reported that cybersecurity informatics specialists do not "get along" well with informaticians specialising in safety and other dependability attributes (Johnson 2016). If there really are such strong social constraints, it does make some sense for guidance documents such as standards to take these constraints into account. On the other hand, software-based system dependability, which includes safety as well as cybersecurity, is recognised as an engineering discipline in itself, and it makes sense that these two of its subdisciplines work as closely together as they can, and also work to overcome social divisions where these exist, rather than encapsulate such social divisions in guidance.

A further consideration is that, in the author's experience, engineers typically like to have as little to do with legal constraints as possible. When Proposals 1 and 2 were being discussed in a small group of engineering specialists, including the author, and the issue of duty of care was raised, the Chair opined that "we are not concerned with the law here" (Anon 2019-2). Taken literally, this must surely be mistaken –

the law in general applies to everything, and must not be ignored. But this statement was intended to convey that conformance with law is a secondary consideration in engineering standards, best left to company legal departments. However, precedent does exist for incorporating legal constraints into engineering. The legal constraint that risks must be reduced "as low as reasonably practicable" (ALARP) was derived in *Edwards v National Coal Board* [1949] 1 All ER 743. It is a principle explained (to some extent) and addressed in (HSE 2001) as well as in an "informative" part of IEC 61508.[6] Various efforts have been made to turn ALARP into an engineering-scientific procedure, but this author regards such efforts, while helpful to engineers who need to understand the principle, to be potentially misleading – the arbiter of ALARP is not some engineering assessor but an English court of law. ALARP is a constraint on safety-critical system engineering but it is not an engineering-scientific procedure.

I would argue that similar considerations which led to the inclusion of engineering processes to address ALARP also lead to engineering processes to address the duty of care. Proposal 2 is far from what the law requires you to do.[7] If it were to prevail, HSE would no longer be able to use conformance to IEC 61508 as a touchstone for appropriate exercise of the duty of care, because that duty requires you to identify and mitigate hazards arising, amongst other things, from cybersecurity vulnerabilities, and IEC 61508 would no longer cover that. So what used to be straightforward, clear guidance conformant with the duty of care would become partial and incomplete. In that case, where else could a system builder and operator go to discharge his or her duty of care completely?

There are various suggestions, all of them with limits. There is a set of standards for cybersecurity in industrial control systems (ICS or IACS), namely IEC 62443 (IEC 62443). There is also nominally guidance on IACS cybersecurity-for-safety (but that guidance is technically poor)(IEC TR 63069, but see Ladkin 2020). The processes recommended in IEC 62443 and IEC TR 63069 may nominally work for systems you can put a fence around, but as a colleague commented at a major national safety engineering conference a couple of years ago, "I work in railway systems and you can't put fences around them" (not in Germany; there are fences in the UK, but they do not keep determined people out) (Braband 2018). There are some sector-specific cybersecurity-for-safety standards; for example, one for IACS, one for nuclear power plants (IEC 62859). Rail has their own and is getting more, automotive (land vehicles) are getting their own, and civil aerospace has it in hand also.

IEC 61508 is designated a "basic safety standard". It applies nominally to everything except medical devices. It surely makes sense to have it reflect the legal duty of care where it can. On identifying and mitigating hazards from any source, it has done so (imperfectly) and this author believes it should continue to do so (while improving the imperfections). It makes much less sense to devolve part of that duty to an uncoordinated number of sector-specific documents. If you are an engineer and your industry sector does not yet happen to have such a specific document, or has a technically poor one, it does not mean you get away with ignoring hazards deriving from cyberinsecurity, because you do not. Best, surely, for safety standards to make that clear.

Peter Bernard Ladkin is a systems-safety specialist with a background in software dependability and logic. His causal accident analysis method Why-Because Analysis (WBA) is used by some 11,000 engineers worldwide. He taught at Bielefeld University and is CEO of tech-transfer companies Causalis Limited and Causalis Ingenieurgesellschaft mbH.

References

Anonymous (2019-1). Personal communication with the author and others concerning Triton/Trisis by a functional-safety manager of an organisation responsible for major industrial plant in March 2019.

Anonymous (2019-2). Personal communication with the author and others, August 2019.

[6] IEC standards are partitioned into "informative" parts and "normative" parts. ALARP cannot be normative internationally because, although it is a principle of English and common law, different principles for risk comparison/reduction prevail in, for example, French law and German law.

[7] And not just English common law. A colleague who chairs the French national standards committee responsible for functional safety and French contributions to IEC 61508 informs me that he has consulted with the French government and been advised that it is also contrary to French law (Ricque 2019).

Bell, Ron (2019). Personal communication with the author (many occasions).

Bowell, Mark (1997/1999). Communication to the author and others on an emailing-list. (unspecified date, 1997-1999).

Braband, Jens (2018). Comments in Workshop 4 of the safe.tech 2018 conference in Munich. Programme available at https://www.tuev-sued.de/uploads/images/1519915089723921992135/safe.tech-2018-programm.pdf.

Health and Safety Executive (2001). Reducing Risks, Protecting People (R2P2). Available from http://www.hse.gov.uk/risk/theory/r2p2.pdf.

International Electrotechnical Commission (2010). IEC 61508, Functional safety of electrical/electronic/programmable electronic safety-related systems, 2nd Edition, 7 parts.

International Electrotechnical Commission (2009-2019) IEC 62443, Industrial-process measurement, control and automation – Network and system security/ also entitled: Security for industrial automation and control systems, many parts.

International Electrotechnical Commission (2019) IEC TR 63069, Industrial-process measurement, control and automation – Framework for functional safety and security.

International Organization for Standardisation/International Electrotechnical Commission (2014). ISO/IEC Guide 51:2014, Safety aspects – Guidelines for their inclusion in standards, 3rd Edition.

International Electrotechnical Commission (2016). IEC 62859 Nuclear power plants Instrumentation and control systems - Requirements for coordinating safety and cybersecurity.

Johnson, Chris (2016). Cyber-Safety: Failures in the use of Cyber-Security Techniques for Safety-Critical Systems and Failures in the use of Safety Techniques in Cyber- Security, Keynote Talk to 24th annual Safety-Critical Systems Symposium, Brighton, UK, 2-4 February, 2016. Programme available at https://scsc.uk/file/378/SSS-2016---v0.71-brochure.pdf.

Ladkin, Peter Bernard (2018). Change Proposal for Subclauses 7.4 and 7.5 of IEC 61508-1, July 2018. Manuscript available to IEC MT 61508-1/2 and IEC MT 61508-2 maintenance team members for IEC 61508.

Ladkin, Peter Bernard (2020). IEC TR 63069, Security Environments and Security-Risk Analysis, Proceedings of the 28th Safety-Critical Systems Symposium, York, 11 – 13 February 2020. Available via https://scsc.uk/publications from February 2020. Shorter preliminary version available at https://www.researchgate.net/publication/333479542_IEC_TR_63069_Security_Environments_and_Security_Risk_Analysis.

Langner, Ralph (2013). To Kill a Centrifuge: A Technical Analysis of What Stuxnet's Creators Tried to Achieve. The Langner Group. Available at https://www.langner.com/wp-content/uploads/2017/03/to-kill-a-centrifuge.pdf.

PAS Global (2019) Triton/Trisis Fact Sheet. Available at https://cyber.pas.com/CyberIntegrity/media/Assets/Fact-Sheet-Cyber-Integrity-Triton-Trisis-Attack.pdf.

Pilz GmbH & Co. KG (2019) Cyberangriff auf Pilz GmbH & Co. KG (in German and English). Available at https://www.pilz.com/message.html, accessed 2019-11-24.

Ricque, Bertrand (2019) Personal communication with the author (many times).

Thomas, Martyn (2019). Personal communication with the author on 2019-07-06.