

IF IT IS NOT SECURE, IT IS NOT SAFE

Why and how safety and security must work together

MIKE STJOHN-GREEN, CEng FIET,
Technical Director
Method Cyber Security Ltd
UK
Email: mike.stjohn-green@methodcysec.com

Abstract

The first reported cyber-attack on a Safety Instrumented System demonstrates that systems important to safety need cyber security measures to avoid their safety arguments being invalidated. But there is a broader justification: cyber security risks arise as a direct result of the nature of networked digital technology, which renders existing safety analysis inadequate to mitigate those risks. Existing standards recognise that safety and security practices need to work together but detailed, procedural best practice is not yet mature. The paper reflects on why cyber security is such a challenge, why safety analysis does not automatically cover cyber security and what current guidance does tell us. It then identifies some specific areas where industry thought-leaders could share how they are dealing with this topic and the paper finishes with a question about the ethics of using of the most advanced networked digital technologies in systems that are important to safety. In considering these topics, the paper draws on experiences in non-nuclear industries, where networked digital technology is being adopted for pressing commercial reasons.

1. INTRODUCTION

“If it not secure, it is not safe” – this is a quote [1] from Professor Robin Bloomfield and his team at Adelard, published in 2013. This message has been around for at least 6 years, but its significance is still not fully appreciated, and how to implement secure safety systems is still being debated [2]. The paper will review progress in developing standards and guidance for the interaction between cyber security and safety and suggest where further work is needed to share emerging best practice.

It is self-evident that safety systems should be protected against the intentional acts of people with malicious intent. Irrespective of their motivation or the means by which they may try to achieve their ends, when designing safety systems, the security of those systems should be considered.

Let us imagine a disgruntled employee who wishes to cause a major accident by causing a storage tank to overflow. In order to do this, they would need to compromise both the control and safety systems. For someone familiar with control and instrumentation components, this would be a relatively simple task, should this person be able to access them.

Physical and personnel security is in place to prevent unauthorised access to the instrumentation. But in today’s world, we must also consider the security of safety systems from those who might exploit networked digital technology to cause serious damage, without exposing themselves to harm and without having to be present on the site at the time. Normal physical and personnel security measures are inadequate alone and cyber security measures will be needed to protect the control and safety systems.

2. IS AN ATTACK ON A SAFETY SYSTEM REALISTIC?

It is clear from what is reported in the press [3] that some malicious actors are exploring the weaknesses in critical national infrastructures. They are developing capabilities to interfere with that infrastructure, should future political, diplomatic or military considerations prompt it. Inevitably, this topic is under-reported – either because the capability has been covertly introduced and has not been detected or because it is not in the interests of the malicious actor or the victim to publicise such attacks when they are detected. Reports [4] have emerged, as part of a campaign to name and shame.

There is one well-documented story of an attack against a safety system – a Schneider Electric Triconex Safety Instrumented System [5]. The nature of the attack suggests that the adversary made a mistake and the

Triconex triple redundancy integrity-checking was triggered, revealing the presence of the intruder. This was reported to be a sophisticated and stealthy attack that required a considerable investment in time and skill. An adversary is only likely to make such an investment to undermine an SIS as a means to a wider goal, such as to cause the process being protected to suffer a major accident.

3. WHY IS CYBER SECURITY SUCH A CHALLENGE?

Since the invention of integrated circuits in the late 1960s and their widespread use from the 1970s, digital technology has replaced electromechanical e.g. relay-based and other e.g. pneumatic-based control systems. The integrated circuits that underpin digital technology have grown in sophistication at an exponential rate, doubling in processing power every 18 months [6]. That processing power has enabled software-intensive systems to be built with ever larger and more complicated programmes. Since the creation of the internet in the 1990s, the digital technology itself is largely designed, created and operated through other digital technology with connections to the internet. This networked digital technology is being used in most of the control systems of today, under the umbrella term of OT.

3.1. Complicated digital technology leads to hidden complexity

New digital technology is built on top of existing digital technology. Rarely is anything built from individual transistors – the electronic equivalent of building it from first principles. The digital technology used to control industrial plant is likely to be built from hundreds or thousands of integrated circuits to form programmable machines that execute software. Integrated circuits may contain millions or in some cases billions of transistors and are the product of high-level abstract designs that are interpreted by other machines, to translate those designs into transistor patterns in silicon.

The authors of those programmes rarely build anything from first principles either. The programmes that these machines execute are generally written in a high-level language and then compiled or translated by other programmes (e.g. compilers) to form the actual executable code that will command the individual transistors in the programmable integrated circuits. These programmes invoke other programmes, or fragments, provided as libraries of helpful routines.

3.2. We deal with complicated and complex systems using abstraction

This technology has become so sophisticated that rarely can one person grasp the detail from top to bottom – from the industrial function being performed to the interaction of the individual transistors that determine whether a valve opens or closes. Contrast this with a panel of relays performing simple ladder logic, where one could follow the wires to determine the control logic. Designers cope with this sophistication by using abstract models for the building blocks that they are assembling:

- the control engineer producing a control systems based on the PLC will regard the PLC and its software as a building block and may use many PLCs, each for a different function,
- the software designer producing PLC software to run on the PLC hardware will regard the PLC hardware, the compiler and software libraries as building blocks,
- the hardware engineer producing PLC hardware will regard its integrated circuits as building blocks,
- the integrated circuit design engineer will regard will regard the integrated circuit design system and software libraries as building blocks,
- and so forth.

3.3. Models used in abstraction can create errors and hide vulnerabilities

At each level, there will be flaws in the building blocks due to errors in their implementation, but those flaws are generally not known to the higher-level system designers (see [7] for the report on Susquehanna Unit 2 nuclear power plant for an illustration). These flaws may create vulnerabilities in the operation of the system. A simplified description of a building block, e.g. describing inputs and outputs but not the complexity within, is vital for complex system design. At each layer of design, it is too easy to assume that the building blocks are perfect. For example, libraries of software functions may have security vulnerabilities that have remained

undiscovered simply because those library functions are unlikely to be fully analysed except by the most rigorous and very expensive assessment. Further, it is impractical to perform sufficient testing to identify all flaws in complicated, software-intensive systems.

There is a further source of flaws, which arise from errors in requirements for those building blocks or assumptions in how they may be assembled to form systems. A simple example is buffer overflow, which occurs when more data than expected is sent to into a computer programme, and neither the sender nor receiver of the data checks the quantity. The effect is that the unexpected data ‘overflows the input buffer’ and overwrites (and corrupts) an unrelated part of memory. Figure 1 is an illustration of a memory map [8] of a simple microprocessor. Consider if the data is destined for the middle of the initialized data area. This memory will very likely have been precisely allocated to many different variables in the programme. The errant data will overwrite the other variables in a manner unintended by the designer.

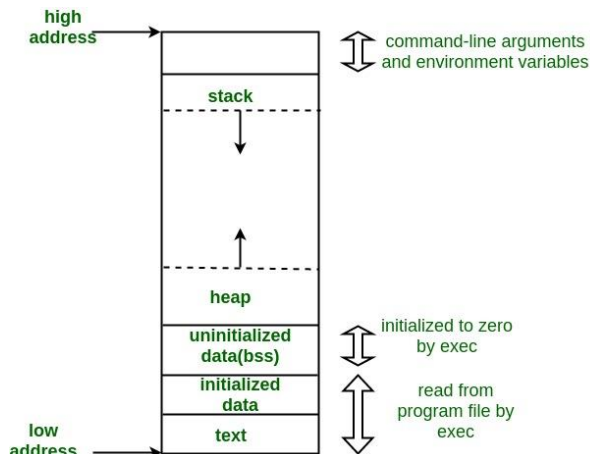


Figure 1 A simple model of a processor memory layout [8]

This can occur because of a mismatch in the specification between the designer of the building block and the system designer who uses the building block. This type of error in requirements results in a vulnerability that remain hidden to the system designer but may be discovered by an adversary and exploited.

4. WHY DOESN'T EXISTING SAFETY ANALYSIS AUTOMATICALLY COVER CYBER SECURITY?

An industrial control and safety system is designed to operate deterministically (i.e. it responds in a particular way to a particular set of input conditions). It is accepted that a mathematical reliability model can be used to calculate the probabilities of various important failure rates. There are some assumptions behind this approach: that the system under control and its control system do not materially change over time; component failures are in general understood and conform to the laws of physics, independent and random (and can be modelled based on these assumptions) and that the humans interacting with the control system are honest and well-intentioned.

4.1. Adversaries may modify our systems in ways we did not anticipate

An adversary is not honest or well-intentioned and may choose to modify the control and safety system from what was originally designed. They may use the control system's own capabilities in a manner unintended by the designer. It would be very hard to model all such possibilities: the safety analysis would have to represent multiple simultaneous failures of the control system or of the system under control, e.g. pumps being commanded to set up oscillations while sensors were also misreading. The design of the attack is unpredictable and unlikely to be considered in a conventional safety analysis. Further, the control system could be ‘caught in the crossfire’ and be an unintended victim of an attack directed elsewhere – such as a ransomware encryption of computer media, where there the intent of the attack is to infect personal computers, not industrial control systems.

4.2. Adversary action is difficult to model with a high level of confidence

Effective safety analysis is based on having sufficient knowledge of component failure rates and modes e.g. a value of the ‘mean time to fail’ for each, based on historical data for the same or similar components. Also, information about design weaknesses, such as those identified from accident investigations, will be available to inform safety calculations. Cyber security does not have the same sharing culture for good reasons – knowledge of vulnerabilities cannot easily be kept from those who may use that knowledge for malicious intent. Consequently, flaws in digital technology (e.g. software bugs) are frequently kept secret until the vendor can provide remediation and users can update their systems. This means that while safety hazards do not change very quickly, the security risks arising from adversary action can change very rapidly.

As has been described earlier, making control systems flexible and programmable also makes them vulnerable to attack. Also, the way systems are built, from building blocks comprising smaller building blocks, means the system designer works with a simplified, abstract model. It may be impractical to model all the vulnerabilities that a system has – even if they were all known at the time of the analysis. And the vulnerabilities will change as the system is developed by the end user, but also as it is “patched” and new features added by the vendors’ software upgrades. These changes may happen very frequently – for example with Microsoft’s weekly updates, referred to by many as ‘Patch Tuesday’ [9].

An attack may be designed so that it “adds” capability over and above what the system was designed to do. For example, an adversary plugging a commercially-available USB device into a system can download a new and malicious programme or introduce a new and uncontrolled wireless connection to the outside world.

Consequently, the cyber security equivalent of a fault tree – an attack tree – cannot be demonstrated to be complete. It can be a useful tool, but it should not be thought of as equivalent to a (safety) fault tree. The attack tree is only as good as the imagination of those who created it and their understanding of the vulnerabilities in the fine detail of the system. They may not be able to imagine all the different ways their adversaries may attack, including all future possibilities as new techniques evolve to exploit vulnerabilities in the system.

4.3. Adversaries exploit high-speed digital networks

Networking of digital technology adds a further challenge. The adversary may attack multiple aspects of a networked control system. Control systems that are connected by fast, flexible networks cannot be assessed in isolation unless the nature of the network connections is highly constrained (to prohibit propagation of malicious activity across the network). Thus, the possible failure modes of individual control systems cannot be analysed by assuming that all its input connections are as they were designed to be – as the attack may have introduced or removed connections in unimagined ways.

Therefore, the paper shows that the model and underlying assumptions currently used for the safety analysis cannot model the actions of a human adversary acting on networked digital technology used in OT. Consequently, the safety analysis that results will not cover cyber security of OT.

5. WHAT DOES CURRENT GUIDANCE TELL US?

5.1. Some standards recognise the need but do not offer detailed procedures

Functional safety standard IEC 61511-1:2016 [10] now calls for a security assessment to be performed, as an addition to the existing hazard and risk analysis. However, the nature of the cyber security risk assessment is not defined and, more importantly, the scheduling of such an assessment is not prescribed. As set out above, the exposure to risk arising from the use of networked digital technology may change suddenly, when a new vulnerability is discovered which can be exploited by adversaries. This means that the cyber security risk assessment should not be an activity that is scheduled by the calendar: the organisation needs to have the competence to recognise when its assessment of risk is no longer valid and revisit the adequacy of its cyber security controls. The cyber security risk assessment must be a continuous, iterative activity performed within a cyber security management system.

A nuclear safety and security standard IEC 62859: 2016 is devoted to the relationship between safety and security for instrumentation and control system in nuclear power plants, though its guidance is more widely

applicable. The standard describes how to integrate cyber security and safety in the design and construction of OT systems, to manage conflicts between safety and cyber security, and maximise the potential synergies between them. It contains a list of examples of potential conflicts between cyber security requirements and safety requirements, as a guide to readers. IEC 62869 does not extend to the level of detail to provide procedures for cyber security or safety experts to reflect this close relationship. Also, IEC 62859 could be unhelpfully interpreted to mean that security objectives should be considered only after the architecture to meet safety objectives has been agreed. Other standards and guidance, e.g. IAEA's NSS-33T [11], indicate that safety and security requirements for nuclear instrumentation and control systems need to be considered together, iteratively.

Cyber security standards for IT environments, e.g. ISO 27000 [12], are mature and have the following common features: a continuous risk management cycle and a catalogue of cyber security controls from which to choose. Those standards are applicable in general terms to OT environments but can deliver the wrong answers because, while IT security considers risk to the information, OT cyber security needs to consider risk to the industrial process of the system under control. Further, OT environments will generally have more stringent requirements for system response times and network latency performance than corresponding IT environments. Some cyber security controls that are commonly used in IT environments may introduce response delays or increases in processing loads that would be unacceptable to the OT's real-time operation.

5.2. Other standards remain silent on the safety-security relationship

Leading cyber security standards for OT environments, IEC 62443 [13], the NIST Cyber Security Framework [14] and – within the UK – the NCSC Cyber Assessment Framework [15], while suited to OT, do not provide guidance on how to integrate safety and cyber security requirements at the time of writing the paper.

5.3. Academics have proposed a systems-based approach

Academics have proposed [16] analysis techniques based on systems theory to provide an integrated approach to both security and safety, but this shifts the focus of the security analysis away from risks arising from specific vulnerabilities to studying the mechanisms that prevent the system from entering a vulnerable system state. In simple terms, this top-down analysis technique focuses on building the systems better in the first place, rather than reacting to vulnerabilities as they are discovered. Arguably, this approach needs to be complemented by a bottom-up approach that addresses the vulnerabilities that were not caught by the top-down approach, as those vulnerabilities are discovered during the operational lifetime of the system.

In conclusion, some safety standards are recognising the need for safety and cyber security to be considered together in the system engineering design and throughout the operational lifecycles. Some safety standards offer general guidance. Cyber security standards are further behind. Crucially, there are currently no codes of practice or detailed procedures for engineers to follow, to tell them what to do differently where networked digital technology is used with systems important to safety.

6. WHAT IS TO BE DONE? - THE ORGANISATION'S RESPONSE

6.1. Continuous cyber security risk management

Organisations need to recognise and act on the issue. In most organisations, there is good level of awareness of the need to maintain the computer security of business Information Technology (IT) assets. However, there has been limited awareness of the impact of poor cyber security on OT, particularly as a result of the use of modern networked digital technology, potentially leaving their OT exposed to cyber-attacks. That awareness is now improving and safety regulators are now demanding the management of cyber security risks where these have a potential to result in harm.

Senior risk owners should expect the nature of cyber security risk to change and should therefore maintain the organisation's capability to track the changing risks arising from networked digital technology change, in order to respond promptly. This requires a cyber security management system, as described in

existing standards. It is likely to call for new ways of working within and between departments, for example calling on the IT security team to feed data that may trigger regular and at times urgent re-assessment of risks in the OT environment.

6.2. Governance: risk ownership, organisational responsibilities

Safety and security communities are often in separate parts of an organisation the objectives of the two communities can be in conflict. For example, the security objectives may require communications to be encrypted to ensure that the end-to-end security of a command and control path is maintained, whilst the safety team may consider this to be introducing another function that could fail, leading to unavailability of the safety function. In simple terms, security seeks to fail-secure while safety seeks to fail-safe, yet an insecure solution may also not be safe for the reasons set out in the paper.

Both the safety community and security community are familiar with making design trade-offs (e.g. on cost, on system usability) to arrive at a practicable system design. A security/safety trade should be seen in the same way as a security/security trade or safety/safety trade. 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.

This in turn raises questions about governance, about who owns the risk arising from the use of networked digital technology in OT environments, who operates the cyber security risk management system for OT and how it interacts with the safety management system. In many organisations, safety and security have no common governance until they reach the company board, by which time it is far too late to make the detailed design and implementation trades to make systems both safe and secure.

The answer to these challenges will depend on the existing management structure of the organisation but it would be valuable to see how best practice is developing.

6.3. Amend procedures and develop new ones – the existing ones are unlikely to be satisfactory

The paper has identified the generalities: safety and security requirements need to be resolved in the system engineering design and throughout the operational lifecycles; and that current standards do not provide detailed guidance on how to do so. It follows that current procedures in many organisations are inadequate to deal with this issue. Standards emerge from best practice that is developed by the thought-leaders in the field, by practitioners who are having to solve these problems then sharing their experience. Areas for such development include identifying:

- how to describe specific requirements for the safe operation of systems in the cyber security risk assessment methods used for OT;
- how to include sufficient analysis of cyber security vulnerabilities in Layers of Protection Analysis, because the assumed independence of those layers of protection may be undermined by cyber security vulnerabilities and the actions of adversaries to exploit them;
- how to include considerations of cyber security in the assessment of hazardous conditions as performed in a Hazard and Risk Analysis. It is insufficient to take an H&RA report and then ensure that cyber security vulnerabilities cannot enable an initiating event. For example, the H&RA report needs to be extended to consider if an adversary can exploit vulnerabilities so as to make realistic a hazardous condition that was previously considered not credible.

6.4. Competence: create multi-disciplinary teams

Unless and until detailed procedures are published, engineers will have to use their judgement and experience, based on their skills and knowledge. However, rarely will all the knowledge and skill to span safety and cyber security for OT be found in one individual. A team approach is almost always required, and that team needs to have access to the knowledge and skills generally only found in IT security teams, of the latest vulnerabilities identified in commodity IT components, such as Windows or Linux operating systems, network protocols and common software tools. The team will also need to understand the industrial processes of the systems under control and the argumentation behind their safety cases.

6.5. Culture: improve mutual understanding

There is a cultural aspect to this: security and safety engineers will typically approach a system assessment from different perspectives. Security engineers are more predisposed to think in terms of what an adversary might do to facilitate an attack, including what may appear to a safety engineer as irrational. However, the security engineers cannot substantiate their arguments with any numerical evidence that would enable the safety engineer to estimate the likelihood of a cyber-attack with any confidence. Safety engineers are more predisposed to assume that the system can be modelled deterministically even when subject to a cyber-attack, despite the system relying on networked digital technology that is built on a very extensive foundation of unprovable software, possibly modified by the adversary.

Whilst safety and security have a common objective of seeking to avoid harm (in a general sense), there are frequently issues caused by the lack of common language and frame of reference. Even where the same word is used by both communities, it may be defined or interpreted differently, leading to confusion.

Therefore, multi-disciplinary team-working may mean including cyber security expertise in HAZOP meetings and when considering layers of independent protection for safety. It also means including safety expertise when considering the defensive cyber security architecture for an OT environment.

Language and culture differ between security and safety. Mutual knowledge and understanding may currently be limited, so organisations may have to take active steps to remediate. Best practice may include job shadowing, cross-training, etc.

7. WHAT IS TO BE DONE? - THE ENGINEER'S RESPONSE

The organisation will have the primary legal and regulatory responsibility to put in place the necessary governance, procedures, culture and competence. However, while those remain immature and incomplete in regard to how safety and security need to work together, the professional engineer has a greater responsibility. In general, the engineer is uniquely placed to understand the risks and act accordingly.

7.1. The professional engineer's responsibility: be cautious and ask questions

The first, key message is one of awareness across the industry: professional engineers should be cautious and ask questions about the risks arising from networked digital technology in systems important to safety. Professional engineers should recognise the immaturity of the safety and cyber security standards where the two disciplines interact. For example, there are claims made by some vendors that safety instrumented systems can be integrated with the control systems that they protect, based on arguments that such combinations still provide adequate safety independence. The evidence and argumentation of such a claim must necessarily make assumptions about the nature of unknown vulnerabilities in a combined control and safety system, and the ability of unknown adversaries to exploit those vulnerabilities now and in the future. It may be that such combined control and safety systems may not be sufficiently secure, even if they are considered to be sufficiently independent to be safe.

There are recommended architectures that show Safety Instrumented Systems (SIS) being connected to the same network as the control systems that they protect. However, since the recent attack on an SIS cited earlier in the paper, at least one vendor has reissued guidance [17] advising that the SIS is not left permanently connected to the control network.

In summary, not everything that can be done with networked digital technology, e.g. putting SIS online, should be done and professional engineers have a responsibility to raise concerns where designs appear to ignore risks arising from networked digital technology.

7.2. Adopt good systems engineering principles to bring coherence to safety and security and resilience

There is a particularly seductive but unrealistic approach that demands that security measures are applied as a separate protective layer outside an existing system – for example one that has already been certified to be safe. Consider the following narrative to illustrate this issue: we have a system that is certified to be safe and we now wish to connect it to a digital network for business reasons so we demand a security mechanism to protect

the safe system from that connection. Security approaches generally assume that no protection is perfect and a proportion of the security resources should be devoted to detection and recovery of the system once the attack has penetrated the protective measures. This demands that the system in our example is modified so that it has the means to detect and recover from any penetration of its defences. By doing so, the system is changed, violating the seductive but unrealistic requirement placed on the provision of its security.

The safety and security goals for a system should be considered together in every part of the system lifecycle to ensure that potential synergies are identified and potential conflicts are resolved. That lifecycle must be defined to include systems states in which an attack has been partially successful.

It is important that safety and security are addressed using tailored views of the same system architecture, that risk analysis considers the combination safety and security goals, not each considered and treated separately, as part of a resilient design. The organisation should define more agile engineering and business processes, as part of its management systems, which enable separate security and safety disciplines to coordinate their activities within a systems engineering framework. Those designs may require modification to respond to the effects of newly-discovered vulnerabilities, newly identified adversaries and cyber-attacks that are discovered while the system is operational.

7.3. Take great care how you define your system boundary

It is important that the system of interest is defined and all interfaces that cross that boundary are well understood. A system boundary used for safety analysis should include those involved in operation and maintenance of the system. This boundary should now reflect the use of networked digital technology and include the supply chain, e.g. where operational maintenance may be remote and cloud-based services may be used. The nature of security risks means that the boundary of the system subject to analysis, which historically has excluded malicious action, may have to be expanded to include the assumed adversaries on the system, whether targeted or incidental, whether ‘insiders’ or external undisclosed third-parties. It also has to consider modifications to the system, such that the system boundary may change e.g. by the introduction of new connections or equipment.

When attempting to analyse the safety effects and likelihood of a system in light of a security attack, the system boundary becomes practically unbounded e.g. where the attack involves changing the configuration of the system and potentially supplementing the system with unauthorised 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 technologies and identifying convenient points at which digital communications can be closely scrutinised and controlled. This will form a more effective boundary that is simpler to manage, even if not a perfect one.

7.4. Make your systems more observable

Simpler systems are easier to make safe and make secure. A corollary is that for a given degree of complexity, the more observable a system, the easier it is to assure its safe and secure operation. Engineers could do more to design systems that are intrinsically more observable in their operation, in order to identify anomalous conditions that may indicate malicious activity or a fault condition. That observability is best designed into the system as a characteristic, although monitoring points on networks can offer the means to make existing systems more observable. Monitoring is only as effective as the operational staff performing the supervisory role and they need to know what constitutes abnormal behaviour of the OT as an indicator of a potential cyber security problem – a less well-defined question than whether the system is entering a hazardous state. Monitoring will increase costs, which will only be endorsed if it is understood as an effective way to reduce recognised risks to the use of networked digital technology in OT.

8. AN ETHICAL QUESTION

The big unanswered question is whether the tools and techniques currently at the disposal of engineers and enforced by organisations are adequate to make the most advanced networked digital technology safe and secure. Given that those tools and techniques inevitably lag behind advances in technology, this means that there

may be some advanced technology that should not be used for systems important to safety simply because best practice does not yet tell us how to make those systems safe enough and secure enough. Organisations and professional engineers all have an ethical responsibility to declare when a design is beyond what can be assured for safety and security, not just build it and then try to fix it. Which of us could stand up in court and say we know where that limit lies?

ACKNOWLEDGEMENTS

The paper is a development of one first delivered in May 2019 at HAZARDS-29, an annual conference of the Institute of Chemical Engineers in the UK that explores major aspects of process safety, particularly in hazardous industries.

REFERENCES

- [1] BLOOMFIELD, PROF, R, STROUD R., Sep 2013, Security-Informed Safety “If it’s not secure, it’s not safe”, MarcOlivier Killijian. Safecom 2013 FastAbstract, Toulouse, France. pp.NC, 2013. <hal00926459>
- [2] BRABAND PROF. DR. JENS, Principal Key Expert Siemens AG, May 2018, Safety and security – principles for railway automation, Signal + Draht Volume 110, Issue 5
- [3] Alert (TA18-074A), March 2018, Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors, Official website of the Department of Homeland Security: <https://www.us-cert.gov/ncas/alerts/TA18-074A>
- [4] COATS, DANIEL R., Director of National Intelligence, Statement for the record for the Senate Select Committee on Intelligence: Worldwide Threat Assessment of the US Intelligence Community 29 January 2019: <https://www.dni.gov/files/ODNI/documents/2019-ATA-SFR---SSCI.pdf>
- [5] MAR-17-352-01 HatMan—Safety System Targeted Malware, 27 February 2019, US DHS NCCICC: <https://ics-cert.us-cert.gov/MAR-17-352-01-HatMan-Safety-System-Targeted-Malware-Update-B>
- [6] A new way to extend Moore’s Law, 8 June 2017, Economist (behind a paywall): <https://www.economist.com/science-and-technology/2017/06/08/a-new-way-to-extend-moores-law>
- [7] U.S Nuclear Regulatory Commission, 4 November 2013, SUSQUEHANNA STEAM ELECTRIC STATION LICENSEE EVENT REPORT 50-388/2012-002-01 UNIT 2 LICENSE NO. NPF -22: <https://www.nrc.gov/docs/ML1330/ML13309A633.pdf>
- [8] Image from Geeks for Geeks, taken 28 November 2019: <https://www.geeksforgeeks.org/memory-layout/>
- [9] Microsoft Patch Tuesday, 9 January 2019, Symantec Corporation: <https://www.symantec.com/blogs/threat-intelligence/microsoft-patch-tuesday-january-2019>
- [10] IEC publications: <https://webstore.iec.ch/>
- [11] IAEA publications: <https://www.iaea.org/publications>
- [12] ISO publications: <https://www.iso.org/>
- [13] ISA-99 standards development: <https://www.isa.org/isa99/>
- [14] US Department of Commerce NIST Cybersecurity Framework v1.1: <https://www.nist.gov/cyberframework>
- [15] UK NCSC Cyber Assessment Framework, October 2018: <https://www.ncsc.gov.uk/blog-post/introducing-cyber-assessment-framework-v20>
- [16] YOUNG, W.E., LEVESON, PROF N., IT, Cambridge, MA, December 2013, Systems thinking for safety and security, ACSAC '13 Proceedings of the 29th Annual Computer Security Applications Conference
- [17] EcoStruxure Triconex Tricon V3, 14 December 2018, Schneider Electric Security Notification: https://download.schneider-electric.com/files?p_enDocType=Technical+leaflet&p_File_Name=SEVD-2017-347-01+Triconex+V3.pdf&p_Doc_Ref=SEVD-2017-347-01