



JRC TECHNICAL REPORTS

Standardisation and Certification of Safety, Security and Privacy in the 'Internet of Things'

*Editor: Gianmarco Baldini
(DG JRC E.3)*

*Authors: Eireann Leverett,
Richard Clayton, Ross Anderson
(Foundation for Information
Policy Research – FIPR)*

2017



This publication is a Technical report by the Joint Research Centre (JRC), the European Commission's science and knowledge service. It aims to provide evidence-based scientific support to the European policymaking process. The scientific output expressed does not imply a policy position of the European Commission. Neither the European Commission nor any person acting on behalf of the Commission is responsible for the use that might be made of this publication.

Contact information

Name: Gianmarco Baldini
Address: Via Primo Maggio 48
Email: gianmarco.baldini@ec.europa.eu
Tel.: +39 0332 78 6618

JRC Science Hub

<https://ec.europa.eu/jrc>

JRC108186

Print	ISBN 978-92-79-77863-6	doi:10.2760/25200
PDF	ISBN 978-92-79-77862-9	doi:10.2760/47559

Luxembourg: Publications Office of the European Union, 2017

© European Union, 2017

Reuse is authorised provided the source is acknowledged. The reuse policy of European Commission documents is regulated by Decision 2011/833/EU (OJ L 330, 14.12.2011, p. 39).

For any use or reproduction of photos or other material that is not under the EU copyright, permission must be sought directly from the copyright holders.

How to cite this report: Eireann Leverett, Richard Clayton, Ross Anderson, Editor G.Baldini, *Standardisation and Certification of Safety, Security and Privacy in the 'Internet of Things'*, European Commission, Brussels, Belgium, 2017, ISBN 978-92-79-77863-6, doi:10.2760/25200, PUBSY No JRC108186.

All images © European Union 2017

Printed in Italy

Foreword

This report was requested by DG JRC E.3 in the context of the Administrative Arrangement Id 34294 between the JRC and DG CNECT to investigate and propose recommendations for an the establishment of a European ICT security certification framework and to assess the feasibility of a European cybersecurity labelling framework.

In this context, DG JRC E.3 (editor Gianmarco Baldini) requested (through contract RTO 261082: Security standardisation for the `Internet of Things') to Eireann Leverett, Richard Clayton, Ross Anderson of the Foundation for Information Policy Research (FIPR) to conduct an analysis on the issues of security certification and the related standards, identify the existing trade-offs and propose policy and standardization options for a potential way forward at European level for security certification.

This report provides the results of the analysis.

Disclaimer: The information and views set out in this study are those of the authors and do not necessarily reflect the official opinion of the European Union. Neither the European Union institutions and bodies nor any person acting on their behalf may be held responsible for any use that may be made of the information contained herein.

Note: This report was submitted to the EU on September 8th 2016, a year before its release for publication. Where the status of regulations discussed in it has changed in the interim, this is noted with the words ‘status in September 2016’.

Executive Summary

As computers become embedded invisibly everywhere, Europe faces significant information security challenges. Emerging problem areas include autonomous vehicles, e-healthcare, smart power grids and smart meters, domestic appliances and even toys; there will be many more. These systems are starting to be known as ‘The Internet of Things’ and they contain new vulnerabilities which can be remotely exploited, with consequent risks to safety and privacy. Many regulators who previously thought only in terms of safety will have to start thinking of security as well. (Indeed, the two concepts are the same in the languages spoken by most Europeans – *sicurezza*, *seguridad*, *sûreté*, *Sicherheit*, *trygghet*...)

Yet the many applications that are acquiring online connectivity and thus exposing their security vulnerabilities to the whole Internet are certified (if at all) under a disparate range of national, industry and other schemes. Insurance underwriters’ laboratories, for example, assess burglar and car alarms, while vehicle safety and building performance are tested by other labs. What happens when we move to smart homes and self-driving cars?

There are several policy objectives we wish to achieve, and available mechanisms include both general provisions, such as the Product Liability Directive, the NIS Directive and the Data Protection Regulation, and the detailed standards and regulations that govern specific industry sectors. However the existing regulators (and standards) mostly take no account of security or privacy threats. Security is complex, and naïve attempts to impose existing security standard frameworks are likely to fail; we give some examples of how they have failed elsewhere.

In this paper we describe the problems and set out some recommendations. The EU needs a multi-stakeholder approach where over-arching regulations on liability, transparency and privacy are coordinated with specific industry regulations on safety and testing. We identify missing institutional resources and suggest a strategy for filling the gap. Above all, the European institutions and regulatory networks need cybersecurity expertise to support safety, privacy, consumer protection and competition. This will be an essential first step towards embedding security thinking in Europe’s many safety regulators.

Contents

1	Introduction	5
2	Historical approaches to regulation	8
2.1	Road transport	9
2.2	Healthcare	12
2.3	Energy sector	16
3	Generic approaches to the problem	20
3.1	Problem statement	20
3.2	Liability as an organising principle	20
3.3	Transparency	24
3.4	The Network and Information Security Directive	25
3.5	Data protection	27
3.6	Attack and vulnerability testing	28
3.7	Licensing engineers, or the firms they work for	32
3.8	Formal methods	33
3.9	The economics of security standards	34
4	Existing security marks and standards	37
4.1	Security standards	37
4.2	The quality and testing standard: ISO/IEC 17025	41
4.3	How security testing and safety testing differ	42
5	Laboratories	44
5.1	CLEFs	44
5.2	SOG-IS	44
5.3	Penetration testers	45
5.4	CITL	45
5.5	ENCS	45
5.6	Underwriters Labs	45
5.7	KEMA	46
5.8	FIRE	46
5.9	Euro NCAP	46
5.10	Summary of European laboratory capabilities	47
5.11	European Standards Organisations	47

6	Requirements	49
6.1	Any EC scheme should build on ISO vulnerability disclosure and coordination standards	49
6.2	Incentivising vendors	50
6.3	Establishing trust and confidence	51
6.4	Collecting and publishing data	52
6.5	A caveat	53
7	Are flaws foreseeable or discoverable?	54
7.1	Novel attacks and accidents	54
7.2	Tool construction and scientific discoverability	55
7.3	Demonstrating harm	56
7.4	Attribution, causation and liability shifting	57
8	Cybersecurity as a public good	60
8.1	Who are the mainstream regulators?	60
8.2	Who investigates safety and security incidents?	62
8.3	Software patches as optimised product recall	63
8.4	Meaningful certification for cross-border trust	64
8.5	Maintenance and evolution of whole systems	64
9	Cyber-covigilance: a new proposal for a multistakeholder approach in Europe	66
9.1	A European Security Engineering Agency	66
9.2	Post-market surveillance	67
9.3	Use of existing labs for post-market studies	68
9.4	Proportionality in causation of harm during investigations and legal remediation	68
9.5	Summary of recommendations	69
10	Illustrative energy sector case study	70
10.1	Real case study	70
10.2	Counterfactual cyber-covigilance case study	71
11	Conclusion	74

1 Introduction

Governments have long had an important role in maximising social welfare by regulating safety and security, where private-sector providers do not have the incentives to do this properly. The motor industry spent many decades competing to decorate cars with chromium rather than fit them with seat belts, until the Product Liability Directive, mandatory safety testing and the provision of crashworthiness information moved them in a more wholesome direction. The regulation of drugs has moved us from the Wild West of nineteenth-century patent medicines to modern standards of safety and efficacy assessed by randomised controlled trials (the safety of medical devices lags somewhat behind). Regulation also plays a key role in consumer confidence; financial regulation and deposit guarantees enable consumers to trust the banking industry despite occasional crises.

The last twenty years have seen the Internet becoming the main vehicle for interpersonal communication and for financial services, as well as a key medium for entertainment, advertising and shopping. The next twenty will see computer-mediated communications embedded invisibly everywhere, from cars and domestic appliances to industrial control systems and medical devices. Large areas of regulation will have to be revisited as the dependability – the safety and security – of computer and communications systems becomes critical to the safety of vehicles, medical devices, and in fact the whole infrastructure on which our societies depend. Indeed, in many languages, ‘safety’ and ‘security’ are the same word (Sicherheit, sûreté, seguridad, sicurezza, trygghet, ...).

At present, the European Union has dozens of regulatory agencies concerned with safety in many different domains. What should this regulatory framework look like a decade from now? Will cybersecurity require a powerful, cross-domain regulator; or will each sector regulator acquire a cell of cybersecurity expertise; or will it be some mixture of general and sectoral approaches; or will we need to develop something else entirely?

The social-welfare goals of a cybersecurity regulator (whether freestanding or sectoral) will typically be some mix of safety and privacy. The former is likely to be dominant in transport while the latter may be more important with personal fitness monitors. Other goals may include national security and law enforcement, competition, the accurate recording of transactions, and dispute resolution. An example where all are in play is smart meters: we do not want the electricity meters in our homes to cause fires, to leak

personal data, to enable a foreign power to hold us to ransom by threatening to turn them off, to allow the utility to cheat us or exploit us through market power, or to make it impossible to resolve disputes fairly.

Achieving these goals will generally depend on mechanisms such as cryptography and access control; assurance that these mechanisms have been correctly implemented and effectively maintained; standards for interoperability; and liability rules that align incentives and prevent actors externalising risks to third parties who are less able to cope with them. It will involve not just writing standards and testing equipment before installation, but monitoring systems for breaches and vulnerabilities, and ordering software updates to deal with both safety and security issues as necessary.

The goals and mission of a cybersecurity regulator may therefore be some mix of the following:

1. Ascertaining and agreeing protection goals
2. Setting standards
3. Certifying standards achievement and enforcing compliance
4. Reducing vulnerabilities
5. Reducing compromises
6. Reducing system externalities

The underlying principle of these individual goals is to maximise social welfare by reducing risk. However, the devil lives in the detail: whose risk should the regulator be reducing – the risk to a dominant industrial player, or to its millions of scattered customers?

The regulators' first task is policy: to determine what needs to be regulated, and why. Yet there may be multiple regulators with different missions: for example, the data protection authorities are concerned with privacy and national electricity regulators with competition. Once the goals are set, these can be elaborated into technical standards in collaboration with industry standards groups and relevant international bodies, building on existing specialist work such as, for example, that done by the US National Institute of Standards and Technology (NIST) in the case of cryptography.

The first goal is self-evident. Without consensus on goals, the mission cannot succeed, as there is no mission. What bad outcomes are we seeking to prevent or to mitigate?

The second goal often entails adapting or evolving existing standards, of which there are already many. There is a huge range of them at all levels in the stack, from algorithms and protocols through software processes to how people should be managed in organisations.

As for the third goal, compliance with standards helps reduce the information asymmetry between vendors and their customers. Business wants to know what it must do in order not to be held liable, and wants predictable tests of compliance.

The focus of the fourth goal is also on reducing the asymmetry between the purchaser and the vendor, but it is dynamic rather than static. Cybersecurity issues are starting to migrate from software products (which are updated monthly to fix bugs and stop attacks) to durable consumer goods such as motor vehicles. Will type approval for cars and medical devices depend in future on regular software updates? A regular update cycle is needed to minimise the amount of time the purchaser is exposed to attacks. Online software updates also greatly decrease the costs of doing product recalls to fix safety problems.

The focus of the fifth goal is largely on reducing the exposure from the viewpoint of an insurer. There are actually a number of defending interests: consumers; vendors; security service companies; computer emergency response teams (CERTs); and finally, government agencies charged with protecting critical infrastructure and other interests.

The sixth goal is once again about reducing asymmetry between vendors and customers. However, its focus is no longer on the vulnerabilities (the technical method of risk activation), but on the overall impact of the externalities. If malware can cause a business to lose money, the regulator might not focus on how to prevent the loss, but on ensuring that the liability falls on the party most able to mitigate the risk. If banks design payment systems then they must carry the risk that they turn out to be insecure; similarly car companies should carry the cost of unsafe software. If autonomous vehicles move liability from negligence to strict product liability, then perhaps they will be sold or leased bundled with insurance, and the safety incentives will be fairly well aligned.

2 Historical approaches to regulation

Let us begin with a quotation from ‘Resilience Engineering’:

“Until recently, the dominant safety paradigm was based on searching for ways in which limited or erratic human performance could degrade an otherwise well designed and ‘safe system’. Techniques from many areas such as reliability engineering and management theory were used to develop ‘demonstrably safe’ systems. The assumption seemed to be that safety, once established, could be maintained by requiring the human performance stayed within prescribed boundaries or norms. Since ‘safe’ systems needed to include mechanisms that guarded against people as unreliable components, understanding how human performance could stray outside these boundaries became important.” [26]

This quote captures an ambivalence about what can be expected of human participants. Safety engineering is both about applications such as transport (where licensed car drivers and airplane pilots can be assumed to have known levels of competence), and also about consumer applications (where products should not harm children or the confused elderly). The same applies to providing security and privacy: the human is not enemy of security, but rather its principal beneficiary. The security engineer’s task is to enable normal users, and even vulnerable users, to enjoy reasonable protection against a capable motivated opponent. How do we protect a confused widow in her eighties against the persuasive fraudster who calls her up and tries to persuade her to install an ‘upgrade’ on her PC, or transfer her savings to a different account for ‘safekeeping’? Human performance can ‘stray’ not just because of error, but because of malicious action by others, and forestalling malice is a much more complex task than making a car crashworthy. Yet as computers and software become embedded everywhere, the regulation of safety will come to include many aspects of information security. It’s not just whether a terrorist can take over my car and use it as a weapon; if a child can use her mobile phone to direct a car to take her to school, what threats do we have to consider in this case?

The task will be to embed security thinking inside standards bodies, enforcement agencies, and testing facilities. Let us examine the history of safety and standards in various contexts.

2.1 Road transport

It took the best part of a century for road vehicle safety to be properly regulated. (Rail safety took much of the previous century.) Car manufacturers initially took the view that if you were injured in an accident you should sue the driver who injured you, and if they blamed the car they should sue the person they bought it from. Attempts to sue car makers for defects started in 1917, but most vendors gave safety a low priority until the campaigner Ralph Nader forced the issue to public attention in the 1960s [5, 36].

His efforts led the US Congress to create the National Highway Traffic Safety Administration in 1966. The US history since then has traced an arc from engineering utopianism, in the form of a belief that crash testing alone would lead to ever-safer cars, to a recognition of the need for real economic incentives by forcing the recall of unsafe vehicles for expensive remediation. The story is told in ‘The Struggle for Auto Safety’ [34, p. 309].

This shift from rules to recalls indicated, as Mashaw and Harfst term it, “a reorientation of auto safety regulation, from science and planning to crime and punishment”.

Even as late as the 1970s, significant battles were still being fought, such as that for the ubiquitous deployment of airbags.

In 1972, General Motors was the industry’s greatest proponent of airbag technology. By 1980, however, GM had become its greatest opponent. The authors credit the auto giant’s defection to NHTSA’s failure to reward innovation and disregard for industry concerns. [34, p. 310]

More recently the insurers have also got involved, with the Insurance Institute for Highway Safety helping develop better crash tests and rating vehicle safety in an attempt to reduce crashes and cut the cost of injuries and property damage.

In the European Union, the most general measure is the Product Liability Directive (85/374/EES) which applies not just to cars but to all manufactured products. It prevents vendors disclaiming liability for injury or death caused by defective products, or damage to the property of individuals (it does not apply to the property of legal persons such as companies).

Can general security principles, regulations and standards also contribute across a range of applications? Yes, they can. Technical standards for common mechanisms such as encryption algorithms (AES, RSA, ECDSA) and security protocols (TLS, Kerberos) are widely used and very valuable. At the policy layer, privacy is dealt with by the Data Protection Directive, and shortly the Data Protection Regulation. Rules for the disclosure of security breaches to affected parties have become embedded in the USA via state-level breach disclosure laws and in the EU by the NIS Directive (though this affects only critical infrastructure, which Member States define differently).

These are complemented by IT industry norms for the coordinated disclosure of vulnerabilities to vendors so they can be patched. The transport industry is struggling to catch up; three years ago, Volkswagen sued the Universities of Birmingham and Nijmegen to delay the publication of a vulnerability in its remote key entry system, which was already being exploited for vehicle theft. It has not however contested more recent similar publications. A regulation covering the disclosure of breaches and vulnerabilities across all industries could do much to encourage laggard industries to catch up with best practice; we will discuss this later.

However general policy measures can't do all the work. Important though the Product Liability Directive was, it was not enough to fix all the issues of safety in transport. Europe has a substantial body of specific regulation on the details of transport safety. Framework Directive 2007/43/EC harmonises the type approval of vehicles; many evolving safety requirements were subsumed under this Directive, such as the additional safety requirements for electric power-train vehicles and battery safety adopted since 2010. Other examples include General Safety (EC) No. 661(2009) covering electronic stability control while the Pedestrian Protocol Regulation (EC) 78(2009) regulates brake assist systems, energy absorbing bonnets and front bumpers. Many regulations are also harmonised internationally through the UNECE framework; the electric vehicle regulations, for example, largely came in via this route.

There are already European regulations governing cybersecurity features of vehicles, such as EU Regulation 165/2014 which makes digital tachographs mandatory for most goods vehicles over 3.5 tonnes; they make tampering harder using smart cards for drivers' licenses and encrypted communications to the sensor.

Regulations have also harmed vehicle security, most notably the Wassenaar Arrangement export controls now enforced via Regulation 428/2009

which limited cryptographic keylength, with the effect that most common vehicle remote key entry systems, having been designed in the 1990s, are now straightforward to crack by brute force – making car theft much easier than it should be.

Now, as we move to autonomous vehicles, the regulation of many aspects of their security and privacy will be embedded in this complex and ever-expanding body of transport regulation rather than becoming separate ‘security’ regulations. A holistic approach is necessary, as the safety of a system depends not just on the vendors and the environment but on the users – and on patterns of behaviour that may have been very deeply embedded. In the USA, we find the same: the NHTSA can regulate not just the carmakers, but also the environment and the drivers (through speed limits).

The move to autonomy will make safety regulation more acute, as vendors will be less able to blame drivers for accidents and the law moves from tort to product liability (which is strict); it will become more complex (along with the software and associated systems); and it will become more dynamic (as software is updated to fix flaws that have caused accidents or security breaches). The regulator’s task will become a lot more challenging.

Similar things can be said about rail safety and air safety. These have long histories and substantial regulatory complexity, which is increasing further with driverless trains and drones of various kinds. However they are beyond the scope of this paper.

The lessons to be learned from our brief survey of road transport are that security in this context is largely an aspect of safety; that while there are some useful over-arching measures to correct generic market failures (such as on liability and transparency) much of the regulatory work will be very detailed and application-specific; and it will evolve constantly over time as vehicles become smarter and more autonomous.

It follows that much of the necessary cybersecurity certification will be embedded in existing testing and certification activities. Thus, for example, when monthly security updates become necessary for cars (as they are for PCs and phones) then it would be natural for the type approval certification under Framework Directive 2007/43/EC to require automatic security update mechanisms, just as it currently requires crash testing and seatbelts.

2.2 Healthcare

Like motor vehicles, medical devices are safety-critical devices, with failures in their design and use being responsible for many deaths.

“Approximately 11% of patients in UK hospitals suffer adverse events, of these half are preventable, and about a third lead to moderate or greater disability or death. Medication errors seem to be one of the most preventable forms of error: more than 17% of medication errors involve miscalculation of doses, incorrect expression of units or incorrect administration rates.” [45]

It is not just the ‘obviously’ safety-critical components, such as infusion pumps and X-ray machines, which cause unnecessary fatalities. The introduction of inappropriate electronic health record systems has also been associated with rising mortality.

“It is tempting, but wrong, to automatically blame the hospital staff [6]. One would imagine that computers, whether embedded in devices or in Health IT systems, would be part of the solution. Unlike calls to improve training or other human processes, if we can improve technology, then everybody can benefit. Yet mortality rates may double when computerized patient record systems are introduced [25]. Healthcare is now widely recognized as turning into an IT problem [27]; computers make every industry more efficient, except healthcare. We clearly need informed, well-founded principles and strategies to improve safety.” [45]

Probably the most famous case of software defects killing people is the series of Therac-25 incidents. A combination of errors led to people being injured in six accidents between 1985 and 1987; three of them died. The combination of events that led to these accidents was rare; the software was complex; and it took some time before a physicist managed to replicate the dangerous behaviour.

“Most accidents are system accidents; that is, they stem from complex interactions between various components and activities. To attribute a single cause to an accident is usually a serious mistake.” [32]

The equipment was dangerous because of poor engineering design, software implementation and testing practices, that are unfortunately all too common in many industrial sectors. Indeed, the Therac investigation had much in common with investigations of security incidents.

“We have tried not to bias our description of the accidents, but it is difficult not to filter unintentionally what is described. Also, we were unable to investigate firsthand or get information about some aspects of the accidents that may be very relevant. For example, detailed information about the manufacturer’s software development, management, and quality control was unavailable. We had to infer most information about these from statements in correspondence or other sources.”[32]

The Therac accidents illustrate the inadequacy of relying on liability laws, as the US victims had to in that case. First, it took five years from the initial incidents to public safety reports; second, it is unacceptable for a victim to bear the burden of proof that a medical device has faulty software.

Medical device regulators should have a rigorous system of equipment evaluation, inspection and certification; accident notification; and recall where necessary. Unfortunately, although the EU has harmonised drug safety regulation in the new European Medicines Agency, medical devices are regulated by Member States, and there is considerable variation. A start has been made via the Medical Devices Directive (93/42/EEC, updated 2007/47/EC) which requires Member States to operate vigilance systems; however devices generally still come into circulation following only a review of documentation and without testing of their functionality and usability (status in September 2016).

Recalls of defective medical devices happen regularly, with several dozen each year, and some of the defects cause death: in the Guidant case, thirteen people died because of short circuits in implantable cardiac devices [49]. An analysis of 23 recalls by the US FDA of defective devices during the first half of 2010, all of which were categorized as “Class I” (meaning a reasonable probability of serious adverse health consequences or death) revealed that at least six of the recalls were likely caused by software defects [42]. Their paper recommends that vendors of regulated medical devices should be required to place copies of their software source code in escrow, to help establish liability after an accident and to allow users to fix the device even if the vendor goes

out of business. However, regulators have to check that the escrowed code is complete, and that it continues to track what the manufacturer ships. Code escrow can form a useful part of the regulatory mix but cannot be seen as a cure-all.

Work by Thimbleby and others has demonstrated that the largest number of avoidable deaths involving medical devices are usability failures. A typical hospital might use infusion pumps from six different vendors all of which have different controls, and these are not always consistent even among the same vendor's products; on some models of the CME Medical Bodyguard 545 infusion pump, for example, the 'increase dose' and 'decrease dose' arrows are the numbers 2 and 0 on the keypad, while on others bearing the same model number they are 5 and 0. This brings to the mind the situation with motor vehicles before WW2, where the arrangement of foot and hand controls varied significantly between models, and someone who could drive a Model T Ford (with its hand throttle and two gear pedals) could not necessarily drive a Peugeot safely, or at all.

The usability failures of medical devices kill about as many people as road traffic accidents do (in the low tens of thousands annually across the EU) yet are difficult for regulators to tackle. Except in egregious cases, such as the Bodyguard 545 just mentioned, each vendor can claim that any incompatibility which leads to user errors is the fault of its competitors [33]. Yet aviation regulators do not permit a pilot qualified on a Boeing 767 to fly an Airbus 340 without a training course; airlines, unlike hospitals, have to internalise the cost of retraining staff for different user interfaces. As a result, Boeing uses the same cockpit design and control layout for the 757 and 767, so pilots qualify on both at the same time. The long-term resolution of the infusion pump safety issue may be along similar lines. But the political pressure does not seem to exist yet.

An even more basic problem is that in many countries (including the USA and the UK) the regulator does not actually test medical devices but instead relies for pre-market certification on a documentation review. The documents required do not include a usability study. So there is no opportunity to come to grips with safety usability, whether directly or indirectly. Despite a series of complaints about safety usability, the FDA appears not to have recalled an infusion pump on usability grounds. As well as the pre-market documentation review, regulation relies on plus post-market surveillance; the latter is not good at rapidly picking up avoidable mishaps that happen one at a time to patients who are mostly very sick anyway.

Cybersecurity attacks are harder to ignore, even if so far they have killed no-one in a hospital. In 2015 the FDA ordered hospitals to stop using the Hospira Symbiq infusion pump, a year after a security researcher showed that the pump could be accessed remotely over WiFi – enabling an attacker to change dosage settings or even use it as a gateway to attack hospital networks. This was a striking response, given the number of patients killed by pump accidents which led to no recalls. However the FDA was unwilling to investigate how many other devices were also vulnerable despite researchers discovering that at least 300 others had similar issues.

The corporate response has been buck-passing, with vendors claiming it is the job of hospital network administrators to block attacks at their firewall, while hospitals claim that the vendors should make devices hackproof. Private action is unlikely to solve this problem; while one large hospital (the US Mayo Clinic) now has its own security requirements for medical devices, few other healthcare providers have the resources to do this [41]. Failing to test security in advance, let alone in the presence of knowledgeable users and medical professionals, leads to trouble.

So the current combination of pre-market documentation review and post-market surveillance is not adequate for safety even for the present world of disconnected devices. A survey of implanted medical device standards and regulation noted:

“European Commission directives do not grant authority to NBs or CAs to require post-approval studies. NBs as part of their review of individual devices can provide guidance for PS, though there is no evidence that studies or registry development are commonly (or even occasionally) required as conditions of approval. Neither the clinical data forming the basis for approved devices nor the existence, if any, of post-approval studies are systematically publicized because there is no requirement for NBs, manufacturers, or CAs to do so.” [31]

By not granting post-approval studies to notified bodies [NBs] and competent authorities [CAs] the EC has effectively eliminated an archive of evidence that would be very useful to security and safety usability researchers alike. This goes somewhat against the spirit of the age of ‘big data’. Of course post-approval studies must preserve privacy, but pharmacovigilance authorities have already worked out how to deal with this.

Luckily, revisions of the Medical Device Directives are underway, which may give an opportunity to improve matters [19]. There is an institutional issue, though: which agency among the European institutions is likely to push for a holistic approach to security and safety during this process?

Safe usability depends on a context, such as that staff are trained to use a device and use it sufficiently frequently to retain this skill. As devices become connected, the network context also starts to matter; if a device can only be operated on a secure network, or if members of the public cannot be allowed within WiFi range, then this should be a condition of use.

We might suggest that the Commission consider modeling adverse event reporting along the lines of the system currently used for pharmacovigilance. This recognises that many things can go wrong with pharmacological supplies after they leave the factory floor; they can be tampered with, stored incorrectly, transported badly, or substituted with counterfeits. Readers from a security background should recognise many of the same issues. The drug licensing system therefore applies pharmacovigilance agreements (PVAs) to specific healthcare providers and patients deemed capable to report on adverse affects. This provides a reporting channel independent of the vendor [53], giving us the data for post-approval studies and a chance of improving safety steadily over time.

There is an issue of institution design, though. The great majority of adverse events reported about medical devices will be safety matters with no cybersecurity element, which may lead device regulators to ignore security concerns. None of the device regulators employ any cybersecurity experts as far as we are aware; in fact, most employ no engineers at all! Most vendors also appear to employ no security engineers, or even any usability engineers, which should be a priority. Safe usability should be the first task of a future European Medical Devices Agency. But cybersecurity might not be far behind.

2.3 Energy sector

ENISA’s own most recent report notes that the energy sector has some of the highest rates of online attacks on critical national infrastructure (CII).

“The most affected CII sectors seem to be financial, ICT and energy.” [48]

It is worth briefly examining the US electric sector regulation eco-system first, as an example of what can go wrong. The USA has five interconnects, nine regional markets, and roughly 1900 bulk power operating organisations. These entities span across generation and transmission, and many more distribution companies above and beyond these 1900. All of these companies have a role to play in securing the grid. The regulators, the North American Electric Reliability Council (NERC) and the Federal Energy Regulatory Commission (FERC), have a small army of auditors to enforce their rules, one of which (NERC CIP) was defined to improve the security of the bulk electrical system; for breaches, NERC can fine companies up to \$1 million per diem. However, their efforts to penalise bad security have produced some unexpected side effects.

“NERC CIP 002 is about ‘Critical Cyber Asset Identification’. Each responsible entity must first identify critical assets and then those cyber assets essential to their operation. Among the critical assets is any generating plant with a ‘black start’ capacity. This means that it can be brought up to power even if the grid is down. In case of large scale blackouts black start generators are used to bootstrap the power grid. Hydro power stations are a good example of plant with an intrinsic black start capability; the operator merely has to turn a valve to allow the water into the turbines, and the plant will spin up. Nuclear power stations on the other hand do not by default have such a capability; they need an external power source to be safely brought up to criticality. In the middle lie fossil-fuel generators, which may or may not have black-start capability depending on whether or not they have auxiliary diesel generators. An alternative black-start strategy is for a plant to have the ability to remain operating at reduced power levels while disconnected from the grid. At the Electric Power 2008 conference, it transpired that plant managers were removing black start capability in order to not have to pay for NERC CIP compliance. This carries a clear cost in terms of system-wide reliability. Some transmission operators were removing IP connectivity from their networks, thereby escaping NERC CIP, while leaving dial-up, Bluetooth and other serial communications into their networks vulnerable. In fact, one of our informants described NERC CIP as ‘a giant exercise in avoidance’!” [2]

The economic incentives were so poorly aligned that NERC CIP ended up compromising the dependability of the US grid (by reducing black-start capability) without materially improving security. Pricing dependability is notoriously difficult; the UK is currently wrestling with possible designs for market mechanisms to provide surplus capacity, of which more will be needed as we move to more variable renewable energy sources. Not all of the components on which we rely for a dependable energy system have explicit prices associated with them. In such an environment, the introduction of new market mechanisms is vulnerable to the law of adverse consequences.

A second lesson is in attitudes to standards. The IT industry is entrepreneurial and freewheeling, with multiple overlapping and competing standards and fairly loose compliance. The electric power industry is different; it has been around since the 1880s and operates expensive equipment that can easily kill. Its engineers are meticulous about complying with every relevant standard, and testing their products rigorously. This leads to conflict when IT bodies offer multiple standards that are not only incompatible, but actually conflict. For example, the Bell-LaPadula model of computer security says that information may only flow up, from a lower level in a system to a higher one, while the Biba model says that it can only flow down. Both are standards; but you have to choose between them.

Another lesson can be learned from smart meter deployment. The UK Government decided in 2009 to deploy smart meters, following Directive 2009/EC/72, with a target of 80% adoption by 2020. In 2010, we queried whether it made sense to fit every home in Britain with a remotely commandable off switch, without making absolutely sure that this could not be exploited maliciously by an opponent [3].

This led to a flurry of activity; later we learned that GCHQ got involved and decided smart meter security had serious concerns [11]. Presentations by their officials made clear that the protection mechanisms they approved focus on preventing large-scale attacks that could let a strategic adversary bring down the grid at a time of tension; they have not concerned themselves much with the smaller-scale problems of whether customers could manipulate the system to steal electricity, or whether the power companies could manipulate it so as to defraud users. This is despite the electricity regulator (Ofgem) being concerned that fraud against customers could get worse.

There has also been controversy in Germany; the German government eventually decided not to install smart meters (on the basis of a cost-benefit analysis) but while this was under consideration the Bundesamt für Sicherheit

in der Informationstechnik (BSI) came out with security standards for in-home equipment that would have locked them down so tight as to prevent any useful interaction between a smart meter and smart home devices such as thermostats and heaters.

The NIS Directive, which we will discuss in detail later, brings in a quite separate regulator. It requires Member States to arrange for firms that are part of the critical national infrastructure to report to some central government agency all security breaches and vulnerabilities. Whether such a body will coordinate effectively or at all with ‘everyday’ industry regulators remains to be seen.

3 Generic approaches to the problem

3.1 Problem statement

At present, cybersecurity regulation in Europe is mostly at the national level, with each Member State vesting its competence in one or more security / intelligence agencies (SIAs) or CERTs, which are coordinated by ENISA. Meanwhile the sectoral regulators for vehicles, health, energy etc are increasingly at the European level but have little to say about the growing use of security mechanisms – and rarely have any cybersecurity expertise.

Where do we want to be in 10–20 years' time? Do we want a single cybersecurity regulator for the EU that covers all sectors and interests? Or do we have cybersecurity policy and technical expertise embedded by sector (e.g., banking, healthcare, energy ...)? Or do we organise it by interest (EDPS to defend privacy, ENISA to forestall external adversaries, another agency to support product safety, yet another to support competition and consumer protection)? Or will it be a matrix of functional and sectoral regulators? If so, will the technical experts be concentrated somewhere, and if so where? What will be the interaction between EU and national regulators?

Do regulators focus on process or outcomes? On standards, on vulnerabilities, on compromises or on liability? Is any of them ever the 'correct' solution, and if not how do we go about deciding which combination to promote in a given context? And are there circumstances in which they are inherently at odds with each other? To what extent can we rely on broad sweeping principles, and when do we need frameworks for much more detailed regulations? Possible general principles include vendor self-certification of security, improved transparency mechanisms, and product liability; we will start off by considering liability in more detail.

3.2 Liability as an organising principle

European law allows people who have been harmed by a defective product to sue the manufacturer. A claimant generally has to establish causation, demonstrate that the harm could have been foreseen, and show that the duty of care was not discharged by the manufacturer. There are significant jurisdictional variations. If the harm is done to the person who bought the goods, then the contract of sale may exclude them from claiming. The EU Product Liability Directive 85/374/EEC was passed to limit this; it provides

that liability for injury or death cannot be excluded by contract, and neither may damage to the property of a physical person (the property of companies is not covered). The Directive also does not cover services.

The two sections of most relevance to safety, security and privacy are Articles 8 and 12. Article 12 states:

The liability of the producer arising from this Directive may not, in relation to the injured person, be limited or excluded by a provision limiting his liability or exempting him from liability.

Article 8 clarifies that actions by third parties cannot erase liability:

Without prejudice to the provisions of national law concerning the right of contribution or recourse, the liability of the producer shall not be reduced when the damage is caused both by a defect in product and by the act or omission of a third party.

Thus an end user license agreement (EULA) cannot invalidate a European user's right to claim damages, and neither does a third party (such as a hacker) tampering with the device. There may also be an action in negligence. For more detail (in a UK context), see Benjamin's Sale of Goods:

- Statutory Liability for Defective Goods: Where defective products cause property damage or personal injuries, liability may also be imposed under the CPA, regardless whether the party suffering such loss was not party to the contract for the sale or transfer of goods. Such liability is imposed upon a class of persons that include the producer of a product, its importer into the EU and those holding themselves out as a producer for damage that is caused wholly or partly by defects in products (s.2(1) CPA).
- Products are regarded as defective where safety of the product is not such as persons generally are entitled to expect, taking into consideration certain statutory criteria (s.4 CPA). The interpretation of that definition has been the subject of detailed consideration, beyond the scope of this article (see, for example *A v National Blood Authority (No.1)* [2001] 3 All E.R. 289). Liability is also imposed under the CPA upon

the immediate supplier of products as well as other suppliers in the chain of distribution where the person suffering damage requests that the supplier identify a party, such as the producer, upon whom primary liability is imposed and that request is not met within a reasonable period (for the full text of this provision, see s.2(3) CPA).

- There is no statutory liability for loss of the product itself and a range of statutory defences is available [10]

However things are not entirely straightforward. A claimant would argue foreseeability of harm by the manufacturer, and certainly, once someone has published a tool that finds software vulnerabilities of a certain type, it would be difficult to claim they were not foreseeable. And if vulnerabilities are found regularly despite previous ones being fixed, then it is foreseeable that more will be found.

The vendor could argue from the exceptions in Article 7 of the Directive, specifically 7(b) “that the defect which caused the damage did not exist at the time when the product was put into circulation by him”, or 7(e) “that the state of scientific or technical knowledge at the time when he put the product into circulation was not such as to enable the existence of the defect to be discovered” or even 7(f) “in the case of a manufacturer of a component, that the defect is attributable to the design of the product in which the component has been fitted or to the instructions given by the manufacturer of the product.” In general, the fact that products work as parts of larger systems is a significant and growing issue for the Directive. It is also a significant limitation that the Directive does not cover commercial property (Article 9).

A further issue is whether the software whose exploitation caused harm is a product (and covered by the Directive) or a service (which at present is not). Firmware in a physical device is very likely covered (though this could be argued; we’re not aware of any case law) while it is quite likely that the software in a server on which an IoT device relies would count as a service. Thus if harm were caused by a defect in (or attack on) GPS navigation, the vendor would be liable for an embedded device like a Garmin Navigator but not for a service such as Google Maps, even though they do exactly the same thing. As we move to a world in which physical devices interact with online services, this needs to be tackled, or vendors will just put the safety-critical functionality in the server to escape liability.

The EU’s own Blue Guide does indeed suggest that it is designed to tackle more than just the initially intended topics:

“Market surveillance should enable unsafe products or products which otherwise do not conform to applicable requirements set out in Union harmonisation legislation to be identified and kept or taken off the market and unscrupulous or even criminal operators punished. It should also act as a powerful deterrent.” [20]

So one of our recommendations is that the EU extend the Product Liability Directive to cover services, and systems that are a mix of products and services. This will become ever more important with the march of virtualisation; we see ‘software as a service’, ‘security as a service’, ‘network as a service’ – all sorts of marketing acronyms ending in ‘AAS’. If the effect of this technological change is to fatally undermine vendor liability for products that kill, then the law must change in step. Thus it is times to carefully consider how to assign liability for life-critical and safety-critical software and services that are composed and intertwined. In fact, it is time for the Directive to be revisited.

Recommendation: The EU should update the Product Liability Directive to cope with systems that incorporate multiple products and services.

A general ability to sue the vendor is necessary, but it is unlikely to be sufficient, as it has not been enough in practice over previous years for vehicles, medical devices and other products with no electronic or service component.

Many Member States have cost-shifting rules whereby the loser in a civil lawsuit pays the winner’s costs; this makes it extremely risky for a private individual to sue a large firm. The USA has tackled this by avoiding cost shifting and making it much easier for claimants to join together in class actions, with the result that lawsuits on issues from car safety to asbestos risks have resulted in real safety improvements. Yet even there, private action has long been accepted as insufficient, and the USA has regulatory bodies like the NHTSA that enact substantial corpora of detailed safety regulation.

3.3 Transparency

Whatever the mix of litigation and regulation, both private claimants and government regulators face a problem of asymmetric information: they know a lot less about a defective product than the vendor does. Not only do they know much less about its design; the vendor also has the accident history, access to which is often the key to both successful claims and effective regulation. Business ethics courses often cite the Ford Pinto case where rear-end collisions had caused the Ford Pinto's fuel tank to rupture, causing fatal fires, and Ford had argued to the NHTSA that the social cost of the burn injuries and fatalities was less than the cost of a recall. After this memo was published, public anger forced the NHTSA to reopen the case, leading to the recall and repair of 1.5 million vehicles in 1977 [51]. In the health sector, pharmacovigilance is the keystone of drug safety regulation, as we have already discussed.

In the world of cybersecurity, transparency is currently provided by security breach disclosure laws and responsible vulnerability disclosure. The first breach disclosure law was introduced in California, and such laws are now on the books in most US states. They require a firm that suffers a computer security breach that compromises personal data to inform the data subjects. Such laws have caused firms to take cybersecurity much more seriously than before; writing to 50 million customers is not cheap, and once there is the prospect of a nine-figure insurance claim, insurers start to take an interest too, as already noted.

Responsible vulnerability disclosure has evolved in the IT industry over the past fifteen years, and represents an emerging consensus on how to deal with vulnerabilities. The researcher or user who discovers a vulnerability might previously have disclosed it publicly, but this places other users at risk and forces the vendor to scramble to roll out a fix; or they might have reported it privately to the vendor, but in this case the vendor might simply have ignored it.

After much experimentation, discussion and economic modelling, the IT industry has stabilised on a system whereby vulnerabilities are reported with a fixed confidentiality period during which the vendor can develop a fix, test it properly, and ship it to users as part of their regular upgrade cycle. This ensures that the vendor has a proper incentive to fix the problem, while minimising their costs. The reporting is often done through a neutral third party, such as a CERT in the case of vulnerabilities exploitable over a network or a

central bank in the case of financial systems; however some vendors operate ‘bug bounty’ programs to encourage researchers to report vulnerabilities to them directly. (There are also brokers who buy vulnerabilities for resale to intelligence agencies and other exploitative users.) There is an ISO standard for vulnerability disclosure, which should inform discussions of this topic [28].

Two of us proposed in 2008 [1] that, as part of the CE process, vendors of devices containing software should be required to self-certify that their products are secure by default, so as to prevent their disclaiming liability for breaches entirely. With network-attached devices this will mean ensuring that vulnerable software can be patched remotely, and it is now time that we spell this out:

Recommendation: Vendors of devices containing software should be required to self-certify, as part of the CE process, that their products are secure by default, and where these devices can be attached to a network they should by default be capable of online software update so that safety and security flaws can be fixed.

The next question is: to whom should breaches and vulnerabilities be reported? This brings us to the NIS Directive.

3.4 The Network and Information Security Directive

The EU’s closest response to US security breach disclosure laws has been the NIS Directive (2013/0027 COD). This directs Member States to require critical infrastructure providers to report security breaches and vulnerabilities to a central agency in each country. This will usually be a security/intelligence agency (SIA) but could for smaller states conceivably be a CERT.

However, Europe has as many definitions of critical national infrastructure as countries. Sweden considers public safety a key element, but aligns its definitions to its ministries. The UK takes a sectoral approach, identifying 13 sectors. As an single example of the complications, one country’s critical infrastructure may reside in another country; Luxembourg depends on generators across the border (many more examples appear regularly in EU cyber-defense discussions). The European Commission itself identified the need for standardisation through the EPCIP back in 2006 [17], but ENISA notes that standard definitions are still an issue [16]. Section 5 in particu-

lar provides methodologies for ‘defining’ critical sectors which some member states have still not finished, and Annex I provides an overview of the different laws in member states that apply.

The NIS directive will create incident reports for incidents around critical infrastructure. However, we still need some work on harmonising the formats of incident reports, and reviewing them periodically, as different data are found to be relevant for safety investigations.

For example, the insurance industry meticulously documents the cost of incidents, but not the technical causes; they want to know if an incident was malicious, but don’t care what vulnerability was exploited. On the other hand intrusion analysts concern themselves with the CVE identifier and CVSS scores of the vulnerabilities used, but don’t count the cost in any detail.

Standardising how we calculate cost is crucial, but hard. Firstly, there are incentives to inflate costs; motives range from increased budgets for cyber defense, through getting police attention, to shifting the blame for incidents to the state.

There is also an important distinction between the public cost imposed on society, and an organisation’s private cost in terms of lost profits. This is the cause of most of the harsh words exchanged between business and government about how much should be spent on cyber defense. A 2011 study showed that the social costs of cybercrime can be two orders of magnitude greater than the amounts actually stolen [4].

Similarly, it will be necessary to standardise the definitions of the security devices that other systems incorporate within their design to achieve relevant security goals. For example, firewalls are often used as building blocks within more complex systems – but they function in many different ways, with different security properties. That’s why buying them can be a difficult experience, they are not as standardised as a catalytic converter, but it would clearly be helpful to have some assurance of what they do. Yet this requires diving into a lot of technical detail about the differences between stateful and non-stateful firewalls, which protocol layer they operate on, and what rules they implement. So quite a lot of work is needed on definitions and low-level technical standards before we can standardise functional testing.

The NIS Directive puts ENISA at the heart of an information-sharing network of national SIAs and perhaps CERTs. However it does not create a resource for doing this kind of standardisation work (much of which is left by default to NIST in the USA). Nor does ENISA appear to have the technical

staff to advise sector regulators, or the incentive to fight for safety to include security as legislation proceeds through European institutions.

What is to happen if (for example) a vulnerability in smart meters will undermine energy-saving objectives, competition policy or users' privacy? It will not be enough to extend reporting: we also have to think about what institutions will analyse the reports and act on them.

First, though, we need to make sure that the data needed by safety regulators to fulfil their security mission are available to them, and not hoarded by ENISA and the Member-State SIAs.

Recommendation: The EU should update the Network and Information Security Directive so that breaches and vulnerabilities are reported to affected users and to relevant regulators.

We will discuss later what sort of agency or agencies should receive those breach and vulnerability reports that are relevant to the regulation of safety, privacy, competition and consumer protection.

3.5 Data protection

The Data Protection Directive, to be replaced in 2018 by the Data Protection Regulation, has for years provided a framework for the protection of personal privacy. Member States have privacy regulators who require that processing of personal information be done according to fair processing principles.

There is established administrative machinery, most notably the European Data Protection Supervisor and the Article 29 Working Party of representatives of national regulators. However this machinery is under strain, for two reasons.

First, the fair-processing rule of thumb of 'consent or anonymise', namely that firms making secondary use of personal data should either get the subjects' consent or redact the data to the extent that it is no longer personal, is coming under strain from Web 2.0, as both consent and anonymisation are rapidly getting less tractable in a world of big data. This will get worse as we move to an Internet of Things, whose sensors will collect ever more personal information: to the location history of your mobile phone will be added your pulse and respiration history from your fitness monitor, and your eyegaze direction history from your augmented-reality goggles.

The second factor is that globalisation is placing the system under ever more strain; as more and more of the systems on which Europeans rely are delivered by external firms (many in the USA) the pressure to relax the regulations is unrelenting.

In the view of the authors, the European institutions should not see this as a problem but an opportunity. The common view in Silicon Valley is that Europe is the world's privacy regulator, as the USA doesn't care and no-one else is big enough to matter. Europe should assume this burden responsibly; if in order to protect our own citizens we have to protect everyone else too, so be it. One way in which the existing institutions can be strengthened is by increasing the cybersecurity expertise available within European institutions. We will return to this point later.

3.6 Attack and vulnerability testing

As well as building on existing laws and norms on liability, transparency and data protection, a fourth generic approach to the problem of assuring the safety and privacy of embedded systems might be mandatory security testing. Perhaps we might simply order all existing regulators to require products they regulate and that contain software to be subjected to attack and vulnerability testing by an independent laboratory as part of the pre-market approval process. This would be no bad thing and has been advocated, for example, in the case of medical devices. However it is not entirely straightforward.

Security and privacy testing is mostly in the hands of private firms who organise penetration testing of client companies and report their findings to the management. While this gives companies an advance view of what they could face at the hands of a capable motivated adversary, it is not as widely used as it might be, as managers are generally loth to pay for bad news that will cause them extra work.

Where a vendor seeks security testing of a product in order to help sell it, there is a clear conflict of interest: he will look for the testing lab that will give him the easiest ride. To tackle this, governments established the Common Criteria evaluation process, under which testing labs are licensed by a national authority. Common Criteria certifications are used for some components of systems in government and banking, but the process is expensive and difficult to apply to whole systems, especially evolving systems; the test results can only speak to vulnerabilities in the device itself in a particular context and usually in a particular configuration. It is generally difficult to

find integration errors that arise when a bug exists between two products, rather than explicitly in one of them. We discussed infusion pumps with incompatible operating instructions; there are many more technical examples, and a whole field of research into ‘security composability’ [8, 13, 46]. It is common for two vendors to each claim the bug is the other’s fault.

This is why penetration testing of whole systems and whole companies is so important. The tester can explore and exploit context around devices, their usage, configuration, and the impact of a given vulnerability, or combinations of them. But most firms lack an incentive to have penetration tests done regularly. There are also methodological complexities.

In most cases, the operational security team of a target company knows the test is going on, which sometimes skews results, as they can often control the test scope, so it doesn’t simulate a realistic opponent. Some firms commission *red teaming*, adapted from the military, where penetration testers are not announced to the firm’s operations staff, and the test is known only to the company’s senior management. The red team can use a wide range of techniques, such as phone-based social engineering, email phishing, or even physical intrusion. This means that the red team is testing not only the hardware, but also the overall capability of the organisations. This event seeks to simulate a real intrusion much more closely and accurately. Unfortunately, red teams usually get in, and firms don’t want to publish this, so such testing cannot easily be relied on by third parties.

Penetration tests, and some equipment tests, are generally covered by Nondisclosure Agreements (NDAs). This has a perverse outcome; that vulnerabilities go undisclosed and often unfixed. As penetration testers move from testing one company to testing another, they find new vulnerabilities. However, they frequently re-use old vulnerabilities when the new company is found to be operating the same equipment.

It is quite common to offer to work with the asset owner to have the vulnerability fixed by the vendor, only to have the asset owner decline. This makes sense given the cost and previous trust relationship between vendor and asset owner; they would rather manage their vendor relationship alone. A few weeks later the penetration tester is in another company using the same vulnerability, with a similar level of success and simulated destruction. Indeed ENISA’s report on testing of Industrial Control System devices noted and highlighted this problem as far back as 2013:

“Consider ways of enforcing vulnerability resolutions Some ex-

perts, especially those from the group ‘Security Test Lab Experts’, noted the importance of keeping some capacity to enforce vulnerability resolution once they have been found and notified. Some experiences around the world have shown that, sometimes, companies do not resolve specific problems, because, for example, they do have economic reasons to do so. This means that some vulnerable systems may stay in production for long periods of time although they have known problems and resolutions, but very closed NDA agreements disable any possible correction enforcement. It has been recommended to keep some level of independence to apply measures. Suggestions for dealing with this vary from applying economic penalties to performing vulnerability disclosures after a reasonable period of time. In any case, it is admitted that any measure would be controversial and could meet resistance.”[15]

This leads to a world where dozens of asset owners may be vulnerable to the same issue, simply because one of them didn’t speak up.

A number of initiatives have been taken to persuade firms to share information about potentially embarrassing vulnerabilities and attacks. The USA set up a number of sectoral Information Sharing and Analysis Centers (ISACs)); the UK has regular sectoral meetings of chief information security officers with officials from the security service; and the EU has the NIS Directive. In theory a tester could also breach an NDA if he considered that a vulnerability amounted to a serious risk to life, but not all would have the courage and such judgement calls are in any case subjective.

We already mentioned the debate among IT vendors and in the security-economics community about disclosing vulnerabilities, which led to an emerging consensus about ten years ago that vulnerabilities should be disclosed responsibly, which means that the tester who finds a flaw reports it to the vendor and also to some responsible body (such as a Computer Emergency Response Team, or CERT, for a network software vulnerability, or to bank regulators for a flaw in banking software). The deal is that the vulnerability will be disclosed and documented after some fixed period of time. This gives the vendor an incentive to develop a patch to fix the flaw, and ship it to their customers in a timely manner. Unfortunately, this practice is not yet widespread in other industries that are starting to rely on software; in 2013, as we noted, Volkswagen took legal action to delay the disclosure of a flaw

in its cars' remote key entry systems.

This was discussed again in a 2008 report two of the authors wrote for the Commission on Security and The Internal Market.

“Publishing quantitative metrics to a wider audience is essential for reducing information asymmetries. We discuss existing statistical indicators, highlighting how they may be improved. We also describe the requirements for constructing comparable indicators. We discuss the options for metrics derived from market price information. Such metrics may be used to differentiate the security levels of software. Another instance of asymmetric information found in the information security market is a lack of data sharing about vulnerabilities and attacks. Companies are hesitant to discuss their weaknesses with competitors even though a coordinated view of attacks could prompt faster mitigation to everyone's benefit. In the USA, this problem has been tackled by information-sharing associations, security-breach disclosure laws and vulnerability markets. There has been discussion of a security-breach disclosure directive in Europe.” [1]

So any standard, certification, or kitemark, should use transparency to incentivise vulnerability reporting from the field, both to get vendors to fix bugs directly, and also as a quality signal so that purchasers can see which vendors have the fewest vulnerabilities, and patch them the quickest.

It follows that while some security standards try to specify a static product in detail (such as the definition of the AES encryption algorithm), the majority have to deal with a moving target. They will typically specify a process whereby products are subject to security testing and fixed when flaws appear.

In the case of automobiles, the current type approval regime will eventually have to incorporate not just pre-market security testing of the vehicle's systems, but a software update mechanism and also a process whereby security patches can be shipped. We expect to see the same with medical devices, electrotechnical equipment and much else.

Standards themselves must be capable of being updated or supplemented in the face of new information on vulnerabilities and attacker behaviour. In the case of encryption algorithms, for example, the discovery of ever more powerful side-channel attacks against AES implementations since the standard was established has forced developers to upgrade them, even though

the base standard remains the same. For example, there is now in many applications a new requirement that the software should execute in constant time.

Second, the notion of adversarial thinking needs to be embedded within the standards process. The whole process must be open to researchers, competitors, suppliers and others – to ensure that the process does not become captured by a small vendor cartel whose real goal becomes liability shifting, and which will use tame testing labs to certify its products whenever it possibly can. Proper adversarial testing, by opponents who are motivated to find vulnerabilities and disclose them, is a powerful antidote. The coordinated vulnerability disclosure norms and programs seen in the software industry show how to organise this and make it work in practice.

We will discuss later how IT industry norms and best practices can be imported into other sectors.

3.7 Licensing engineers, or the firms they work for

In Canada, engineers working on safety-critical systems are required to be licensed professional engineers. To achieve this status, engineering graduates must undergo a programme of on-the-job professional training, and when they qualify they may wear an iron ring, which confers social status. It is often suggested that Europe should move in the same direction.

Engineering culture varies across Member States, with engineers having higher social status in Germany and France than in some others. In the UK, engineers can achieve ‘Chartered Engineer’ status by a process similar to the Canadian one, but most employers ignore this and salaries are not particularly high. The EUR ING title is made available by the European Federation of National Engineering Associations (FEANI) and provides a harmonised qualification certifying at least three years’ engineering education at university followed by at least two years’ practical experience.

Such qualifications do no harm, but they do not necessarily solve the problem in hand. If a medical device vendor has user interfaces designed by the same electronics engineers who design the device’s circuit board, who lack expertise in usability and are unaware of their ignorance, then they can as easily be ignorant of security. All accredited computer science courses contain at least some lectures on cryptography and computer security, but this is not necessarily the case with electronic engineering degrees.

There are also various schemes for accrediting companies, and teams

within them, from ISO 9001 through the CMU SEI Capability Maturity Model to various sectoral schemes, for example in aerospace. Such sectoral schemes might serve as a conduit for getting staff to undergo security courses as part of their continuing professional development. This is all very worthwhile but it is slow and should be seen as something that supports regulatory goals rather than helping achieve them directly.

3.8 Formal methods

There has been an enormous amount of research over the past fifty years into whether it is possible to prove programs correct, and many of the most eminent computer scientists have written on this issue. There are formal specification languages such as Z and Lotos that can be used to express program behaviours that are desired or forbidden; specialist logics such as the Burrows-Abadi-Needham (BAN) logic, used for verifying authentication protocols; special calculi, such as the pi calculus which has been used to verify the protocols underlying 3G communications; and software tools, from the theorem-prover Isabelle which has been used to verify the TLS protocol to Hol which was used to verify the Viper chip.

These tools can be very useful for verifying compact designs such as protocols, or components of a chip. The use of such tools can be fiddly and thus expensive in terms of highly-skilled (PhD grade) labour. But they come into their own where a component is security-critical but may be written by an untrusted party; for example, Microsoft uses the SLAM model checker to verify device drivers that ship with Windows.

An alternative approach is that taken by firms like Coverity, whose static analysis tools are used by an increasing number of companies to find bugs in large, complex programs. The philosophy here is not to find all the bugs (many programs are just too big and complicated for that) but to find very large numbers of bugs efficiently. Thus, for example, when a new type of vulnerability is discovered, security-conscious software developers will not just fix the bug, but fix their tools to find all similar bugs.

The problem with using all formal methods is cost. A firm that starts running Coverity on its code base may find tens of thousands of bugs that never caused any trouble in the past, and have to delay shipment for several months while they are fixed. Once this task is done, the tools will find further bugs as they are written, so there is no great additional cost for a while – until Coverity ships its next version, at which point there is once more a

backlog of bug fixes.

Firms like Microsoft have learned the hard way, over many years, that the use of such tools is worthwhile in the long run, despite the added cost, delay and pain. However many firms are reluctant to join this club, and here a push from a regulator may well be worthwhile. We will discuss later what form this might take; first we must consider what standards might be promoted, and what they achieve.

3.9 The economics of security standards

Although we now know a fair amount about the economics of security breach reporting and vulnerability disclosure, there has been relatively little economic analysis of security standards. Ordinary technical standards have been studied in the context of standards races, patent pools, regulatory capture and innovation generally, and are the subject of a significant literature. But security standards have attracted less attention [38].

The economics of cybercrime are not zero-sum; attackers profit very differently from defenders. A 2011 study of the economics of cybercrime showed that for every \$1 an attacker earns, the defenders are spending from \$1 in banking to \$100 in the case of the more modern cybercrimes [4]. A regulator aiming at socially optimal outcomes must therefore understand the different actors' incentives,

- Attackers often make money in a completely different way from their victims. They may steal money directly, or subvert advertising websites as a distribution engine for malvertising or malware.
- Defenders' losses are sometimes just the attackers' gains, such as in credit card fraud. But much more frequently, they are very different. A user whose PC is infected by malware may be advised by the shop to buy a new one rather than cleaning up their old one; the small gains made by the attacker are dwarfed by the profit made by the shop, the PC maker, and its suppliers.
- Vendors may operate under competition, as an oligopoly, or a monopoly. Especially in the latter cases, response may be slow. The absence of effective product liability for most software may exacerbate this.

- Insurance of cyber risks is a complex issue. Cyber risks are usually covered as part of general business insurance and for years were not a sufficiently large part of total claims for insurers to pay much attention. The spread of security-breach disclosure laws in the USA has changed that; having to write to 50 million customers to notify them of a breach is expensive enough to matter. However reinsurance of such risks raises issues of risk correlation. These complicate any simple analysis based on frequency and severity of attacks, and the effectiveness of mitigating solutions. Insurers do have consultants to make an assessment of major customers but cannot afford to perform penetration testing of every customer. Up until now, their focus has been on the quantum of losses rather than the actual vulnerabilities exploited, as they did not have detailed technical information about most of their customers. Transparency of breach and vulnerability reporting can thus be of real value.
- Societal risks have been demonstrated in incidents such as the December 2015 Ukraine power outage, where nation state actors in particular may have wider systemic impact. This may in extreme cases have measurable effects on the GDP of the affected state or even impact supply chains in a global economy. A second societal risk is loss of confidence in online transactions as a result of the rising tide of cybercrime. This makes people less likely to shop or use government services online, increasing costs for all, and slowing down innovation and growth. Its effects are orders of magnitude greater than the actual sums stolen by cybercriminals directly.

Any future EU security lab, standards body or certification scheme may therefore have multiple goals, including to:

1. drive up the cost for attackers and reduce the income they can generate;
2. reduce the cost of defence and also the impact of security failure;
3. enable insurers to price cyber-risks efficiently;
4. reduce the social cost of cybercrime and social vulnerability to attacks.

The optimal balance is likely to vary from one sector to another, and also to vary over time.

In the next section we will survey some of the current standards which we can use as building blocks, or as lessons about what to avoid.

4 Existing security marks and standards

There are many security standards at all levels from components to whole systems, covering many aspects of product and process. Some are good and others less so; some are necessary in some contexts, but none are sufficient; some standards have been developed with rigorous peer review and testing while others have been pushed through to justify what particular industries were doing already, or even to dump liability for failure on to third parties. In short, security standardisation is a complex and confusing subject. Although standards are far too numerous to be discussed at length in this paper, we present a quick sample of them.

It is relevant to note that standards generally concern themselves with things that are easy to test (such as conformance with a cryptographic or architectural specification) rather than with the harder problem of writing secure software. Most of the exploitable vulnerabilities occur not because someone used a non-standard cryptographic algorithm but because a programmer implemented a standard algorithm but in software that turned out to be exploitable.

4.1 Security standards

- NIST SP 800-183
 - This special publication on “Networks of ‘Things’” is closest to the subject matter of this paper. It sets out a framework for analysing security in IoT in terms of five primitives – sensors, aggregators, communication channels, eUtilities and decision triggers. Its contribution is at the level of architecture and terminology, as an aid to security analysis, rather than anything directly testable. However it gives some idea of the potential complexity and the range of other standards and tests that may be necessary in building and certifying an IoT system.
- NIST Cryptographic Standards
 - The US National Institute of Standards and Technology (NIST) has since the 1970s developed a series of standards for cryptographic algorithms, modes of operation and protocols, starting with the Data Encryption Standard (DES) in the 1970s, followed

thirty years later by the Advanced Encryption Standard (AES) and secure hash algorithms. These are very widely used. Adherence to standards is not however sufficient for security as implementations should usually run in constant time to forestall side-channel attacks, as noted above.

- NIST SP 800-82
 - This is a best practice for the secure deployment of Industrial Control Systems (ICS). It is aimed at deployment, and concerns itself with network architecture and the use of security features such as firewalls for protection. It is not as applicable to programming ICS systems securely.
- IEC 62443
 - Sometimes referred to as ISA99 or ISASecure, this standard covers network security of industrial control systems. It aims to segregate networks to protect against contagion. It does not focus upon how to develop or code securely when producing software intended for critical infrastructures. Most notably, it is not as relevant for the testing of the security and privacy preserving aspects of industrial system devices.
- DO-178B
 - This standard concerns itself with software intended for use in airborne vehicles and the certification of devices. The primary aim is to ensure software is fit for safety critical purposes. Thus it is intended more to protect from error and mischance than from malice.
- IEC61508
 - This standard covers functional safety, is aimed at the electrotechnical industry (though used elsewhere too), and is at a higher level than EN50128. It provides a methodology to assess the risks to systems and determine the safety requirements, and introduces both safety integrity levels and the safety lifecycle. It supports the certification of components for use in safety-critical systems.

However its focus is on bounding failure probabilities, and it does not consider a malicious adversary.

- EN50128
 - This standard does concern itself both with security and safety certification of software, and follows IEC61508. Unfortunately, it is tailored for the rail and transport industry, and is not seen as immediately applicable to other industries. It supports safety integrity levels, and enforces both rigorous quality assurance and safety critical component testing, using independent assessors. It also recommends the use of languages other than C, various kinds of static analysis, and traceability of all code to requirements.
- ISO 26262
 - Another standard that follows on from IEC 61598, this standard is about to the functional safety of road vehicles. It does cover the full lifecycle of development, but does not refer to best practices for security.
- ISO 15408
 - This is the Common Criteria, a framework operated by a network of national authorities, typically the signals intelligence agency in each member country, which certifies a number of labs as Commercial Licensed Evaluation Facilities (CLEFs). A vendor can define a Protection Profile (PP), which specifies the security policy a product should enforce; it can have the PP evaluated by a CLEF and registered for wider use; it can then have a CLEF evaluate whether a product meets a particular PP to a given level of assurance. Evaluations are registered with the national authority, and are then recognised in other participating countries. This standard does not concern itself with safety, except incidentally. It is a slow and expensive process, with an EAL 4 certification typically costing in six figures and taking many months. It is used mainly for testing computer systems for sale to governments. The protection profiles against which these tests are performed are harmonised in the EU by SOG-IS, which we describe briefly in the next section.

- Multiple Independent Levels of Security/Safety (MILS)
 - The MILS specification does concern itself with both safety and security, but applies mainly to operating system design in general. It can be used in combination with Common Criteria techniques, but is more of a design principle than a testable standard. It is good at preventing certain types of vulnerabilities, but is agnostic about others, assuming they are the user’s responsibility.
- Euro-MILS
 - This standard takes into account virtualisation and kernel isolation in much the same way as MILS. It is meant for both safety and security, but is also slow and cumbersome to test against. Primarily useful for high assurance systems like aircraft, defense, and intelligence purposes, but too expensive for IoT devices.
- IEC 62304
 - This standard is devoted to the total software lifecycle of medical devices. It does recognise that software of unknown origin should be vetted with a risk-based approach, and not be used if at all possible, but is a self policed strategy at heart.
- ISO 27001/27002
 - ISO 27001 sets out to ‘provide requirements for establishing, implementing, maintaining and continuously improving an Information Security Management System (ISMS)’ while 27002 has a list of possible controls. Essentially, these documents provide a framework for a large organisation that seeks to measure and evaluate how well it does information security management; they make it susceptible to internal and external audit processes, and are basically seen as audit checklists. However, they are fundamentally about companies securing their own assets and operations, not about making products that protect their customers.

- ISO 9001
 - This standard is relied on by a number of the above standards. It is a general quality assurance standard about the repeatability of a firm’s management processes, with a systems approach to management, factual decision-making and a commitment to continuous improvement.

We also mentioned above that much work needs to be done on mundane matters such as definitions. It was not until a British standard plug (and corresponding electrical socket) was defined that testing for safety purposes really progressed. Indeed in October 1971, the specifications for a ‘British Standard finger’ were published, giving test facilities clear guidance on the size and width of probes for checking access to ‘live’ parts.

Similar foundational standards are needed for a range of cybersecurity components, along with standards for how conformance will be checked. Perhaps the EU is content for NIST to tackle such work, at its own pace. However there is an argument for Europe to have a technically competent agency that can initiate and contribute to standards as European requirements evolve.

4.2 The quality and testing standard: ISO/IEC 17025

This standard perhaps merits discussion in greater detail than those summarised above. It standardises the accreditation and certification of testing (or calibration) laboratories. Essentially it testifies to the technical competence of the laboratory, its staff, processes, and quality assurance. It is globally recognized, and pertains to testing facilities of all types from automotive to zoological, but most crucially safety. It details auditing, information ownership, and (importantly for security!) how to deal with non-conforming test results. To the knowledge of the authors, no security and privacy testing labs have sought or achieved this accreditation. This is something that will very likely change over time, as security and privacy research laboratories seek to get involved in safety and to go beyond CLEF approval. This is likely to aid the debate between security and safety for reasons we will now discuss.

4.3 How security testing and safety testing differ

Security testing differs in a crucial way from the more familiar safety testing with which engineers are familiar. We are looking not for failures that occur under typical operation and loading, but failures that can be caused to happen by a malicious opponent who ingeniously and creatively applies unexpected combinations of operation, loading and context. Vulnerability discovery is thus by its nature exploratory, rather than scripted and deterministic. It is the goal of a decent security tester to ‘think outside the box’, and if need be find ways around any standard.

Some exploits are probabilistic in nature; they may involve ‘fuzzing’ where the tester tries large numbers of possible inputs to try to get an output beyond the safe (or secure) range. It may also be difficult to reproduce security vulnerability results, especially months or years after they are initially reported – as the software may have been patched, or the configuration may have changed, or the environment may be different. Complex systems have many dependencies and large-scale automated testing is essential to discover obscure flaws that manifest themselves only rarely with particular combinations of input and context. This can have significant consequences for the investigation of accidents and the determination of liability, unless steps are taken at the time to reproduce the problem or at least to keep extensive logs and records for later analysis.

In the Therac-25 incident mentioned in section 2.2, the fatal accidents occurred because of the precise timing of user inputs that was needed to expose the flaw, and timing wasn’t considered important until a physicist at a hospital where a fatal accident had happened dedicated the time to experimenting until he could repeat the conditions that led to the accident. If we only test the input conditions that have led to accidents before, we will miss new combinations that might lead to harm; so it is now best practice to test large numbers of random inputs too. Repeatability of tests – the gold standard in safety testing – is no longer enough.

Resilience requires the ability to detect, block and recover from the actions of capable motivated opponents. Detection requires good record keeping plus someone with the incentive to watch, while recovery involves patching systems. In fact, remote patch management can be seen as optimised product recall.

A vendor with a good patching system can replace an insecure piece of equipment with a patched one in a few minutes of download and reboot

time. This is by far preferable to product recalls in the physical world, but will require a large cultural shift for vendors of devices from cars to white goods who see avoiding a recalls as a management priority. Embracing it instead will be a seismic shift.

Indeed, although computer vendors such as Microsoft and Apple now run monthly update and patch cycles, not all mobile phone vendors have embraced this yet. Extending patching to TVs and airconditioners was discussed in our 2008 report [1] but has not happened yet.

The Commission should push by regulation to ensure that network-attached programmable devices can be patched remotely to fix safety and security flaws, as we recommended in section 3.2 above. As for how the patching mechanisms should be managed once they have been deployed, we will discuss that later.

5 Laboratories

A number of laboratories deal with privacy security and safety. In this section we examine a few that have developed standards or testing facilities. These labs fall into roughly one of two categories: great at security or privacy testing, but without experience of developing, amending, or certifying against safety standards; and the reverse, with plenty of safety compliance and conformance testing experience, but no adversarial approach or seasoned security engineers.

Security engineers and safety engineers are two separate tribes. The authors have seen many safety experts dismiss security concerns simply because they did not see it as their job. Standards bodies and testing labs must learn to stop ignoring security concerns as “out of scope”.

5.1 CLEFs

A number of firms in Europe have been certified as Commercial Licensed Evaluation Facilities (CLEFs) under the Common Criteria, and do certification work that is recognised across participating states. They mostly evaluate software products for government use, though there is some work on products such as smartcards that are sold into the banking industry. Such firms can also be favoured for work that falls outside the scheme, which may be represented using a term such as ‘Common Criteria Evaluated’ although this carries no force under the standard.

5.2 SOG-IS

The Senior Officials’ Group – Information Systems Security (SOG-IS) consists of technical experts from the national information security agencies of Austria, Finland, France, Germany, Italy, the Netherlands, Norway, Spain, Sweden and the United Kingdom. They coordinate the protection profiles to be used in Common Criteria evaluations of systems for government use in the EU and EFTA, to ensure mutual recognition of certification. SOG-IS also coordinates protection profiles where the EU launches a Directive that must be transposed to national law and has a relevant information security aspect. Its remit includes smartcards and devices with ‘security boxes’ or cryptoprocessors, thus including things like smart meters, taxi meters and tachographs.

5.3 Penetration testers

A larger number of firms offer vulnerability assessments, penetration testing, and red teaming than ever before. The quality of work is variable, as is the assurance on offer. Some penetration testers work for ‘big four’ accountants, established system houses, or approved defence contractors; others specialise in particular sectors. Some Member States have pushed firms in particular sectors (such as banking) to use certain approved firms, which has led to some pushback from firms worried about possible conflicts of interest.

5.4 CITL

Another approach to vulnerability assessment is for firms to have charismatic leaders. As an example, the Cyber Independent Testing Laboratories is run by Pieter Zatko, aka ‘Mudge’, a well-known security researcher. Their initial stance seems to be using automated tools to grade thousands of computer programs at scale, and provide security metrics.

5.5 ENCS

The European Network for Cyber Security resides in the Netherlands, and has a multi-stakeholder approach. They have tested vehicle-to-vehicle communications, smart grids, and industrial control systems. They have become notable for breaking a few smart grid cryptographic protocols, and work closely with both asset owners and vendors of industrial control system products. They have also designed new protocols, fixed many vulnerabilities, and trained operators in identifying and responding to cyber attacks.

5.6 Underwriters Labs

The Underwriters Labs (UL) in the USA has worked on shock and fire hazards for more than a century. It seems a natural place for security and safety research to become more entwined, and they launched a Cyber Security Certification Program in 2015. The new UL program expects to improve device security, with an eye towards reducing risk for cyber insurers. They have also released a new set of standards devoted specifically to cyber security the UL2900 series. The standards are not free, but the fees are meant to

go towards enhancing in them in a regular operational cycle that matches adversarial operations.

The UL does perform testing of electrical safety to certify products destined for the EU, according to the EU's own standards. It is one of the few labs offering both safety and security testing.

5.7 KEMA

KEMA has been testing and certifying electrical equipment for the transmission and distribution of electricity in Europe for nearly a century. The lab is accredited to the ISO 17025 standard by the Dutch National Accreditation Council and has extensive experience with safety systems testing, particularly in the electric sector. There have been some successful efforts at security testing from KEMA, but primarily their testing is compliance-based.

5.8 FIRE

The Future Internet Research and Experimentation (FIRE) initiative, which was launched by the EU in 2010, promotes a more clean-slate and adversarial approach to producing new products and protocols for the Internet of Things. They offer lab environments for pre-market experimentation, including some random and adversarial network events to stress and improve the resilience of products.

5.9 Euro NCAP

The European New Car Assessment Programme (Euro NCAP) is a European car safety assessment programme. Founded in 1997 by the UK Transport Research Laboratory for the UK, it is now backed by the EU and a number of other Member States. Euro NCAP Facilities across Europe collaborate to offer a star rating system of car safety as a market signal for car buyers, on top of the far more extensive, and costly, type approval work including crash testing [21]. This is an example of an ecosystem into which, in the long run, we have to embed a culture of adversarial security engineering.

5.10 Summary of European laboratory capabilities

Laboratories in Europe tend to study either safety or security. However, safety and compliance tend to be rule-based, while hackers delight in abusing rules. Just as safety progresses with the study of each accident, security progresses with the study of every exploit and every adversary. As safety starts to require security, compliance needs to be supplemented with adversarial thinking. Asking junior analysts to break a system improves its resilience and robustness, as well as honing their skills.

This has implications for the regulatory environment. In the presence of regulatory capture security wanes. This is not only for economic reasons, but also because regulatory capture makes the environment averse to change. Standards bodies run entirely by vendors reflect this perfectly.

Testing labs can serve a vital role in protecting the Internet of Things, but a significant role will also be played by European standards organisations.

5.11 European Standards Organisations

Europe has three main relevant standards organisations:

- The European Committee for Electrotechnical Standardization (CENELEC) is an institution of the EU and EFTA, with access to over 33 countries (most of them EU Member States). It is also diversified across industrial sectors and has been involved in producing safety standards for electrical engineering, electric vehicles, medical equipment, railroads, smart grids, smart metering and much else.
- The European Committee for Standardization (CEN) has 31 national members and its mission is removing trade barriers by standardisation. It operates through a number of technical committees; for example CEN TC 251 sets interoperability standards for healthcare informatics and plays a key role in the development of systems such as electronic prescribing and the electronic reporting of lab test results to doctors' offices.
- The European Telecommunications Standards Institute (ETSI) is a membership organisation of some 800 firms which have worked together to produce standards for the telecommunications industry, the most high-impact of which was GSM, the standard for digital mobile

telephony now used in most countries round the world. GSM, like later mobile communications standards, incorporates cryptography and other security mechanisms.

ETSI is explicitly an industry body, while CEN and CENELEC operate via technical committees most of whose members are from industry. This raises the issue of regulatory capture, of conflict of interest between the vendor community and other actors, and the thorny problem of how we move away from insecure legacy systems.

The major limiting factor of a standard's rate of change is often backwards compatibility with earlier equipment or protocols. This is a legitimate concern, in a world where things need to continue to interoperate. But the implication in the world of control systems is that it may take 25 years to replace existing communications protocols. These protocols evolved in a world of private networks made up of hard-wired connections, leased lines and private radio links; anyone who could connect to a sensor could read it, and anyone who could connect to an actuator could command it.

Some 15 years ago, control systems moved to IP networks for reasons of cost, and only afterwards did operators realise that their systems were wide open to attack. The belated response has been re-perimeterisation. Work is afoot to add authentication to control-system protocols such as modbus or IEC-61850 but it is proceeding slowly, and the need to interoperate with legacy equipment will ensure that unprotected systems are still being installed and extended for years to come. We noted a similar problem with medical devices, whose vendors consider it the hospital's responsibility to run a 'secure' network (i.e. an environment in which vendors need not worry about security, and can ship devices with well-known hard-coded root passwords).

There is a role here for an external functional regulator to give the industry a push. Without it, standards will change by consensus, and the convoy will therefore move at the speed of the slowest ship; often an insecure device or protocol will provide a means of ingress for years before it is remediated. The point here is, simply, that security standards must evolve as attackers do.

Recommendation: The Commission should investigate how to move European standards bodies in the direction of considering safety and security together in future standards work.

6 Requirements

As discussed above, breaches and vulnerabilities that affect critical infrastructure should now be reported to ENISA, but there are no general mechanisms for reporting other flaws or incidents, for reporting to other regulatory or enforcement bodies, or for certifying the safety of affected devices. The authors have reported vulnerabilities in payment systems to the European Central Bank, to the police, and to Visa (which told banks, some of whom eventually pushed vendors into action). Others have reported vulnerabilities in car electronics to the vendor and received legal action from its customer the carmaker.

What then are our requirements for any new system that covers safety and security certification, and the remediation of relevant vulnerabilities?

Here we elaborate a few core requirements.

6.1 Any EC scheme should build on ISO vulnerability disclosure and coordination standards

Any new scheme should respect the work already done in vulnerability management, and encourage stakeholders to work within a coordinated vulnerability disclosure process. In particular, liability rules should not drive vendors to avoid working with external researchers, but rather encourage this.

Luckily, there is a relevant ISO standard that already covers many of the key issues in the vulnerability disclosure process.

“ISO/IEC 29147:2014 gives guidelines for the disclosure of potential vulnerabilities in products and online services. It details the methods a vendor should use to address issues related to vulnerability disclosure. ISO/IEC 29147:2014

- provides guidelines for vendors on how to receive information about potential vulnerabilities in their products or online services,
- provides guidelines for vendors on how to disseminate resolution information about vulnerabilities in their products or online services,

- provides the information items that should be produced through the implementation of a vendor’s vulnerability disclosure process, and
- provides examples of content that should be included in the information items.

ISO/IEC 29147:2014 is applicable to vendors who respond to external reports of vulnerabilities in their products or online services.” [28]

A companion standard, ISO/IEC 30111, on “Information technology – Security techniques – Vulnerability handling processes” provides the structure for verifying a reported vulnerability, whether discovered internally or externally, developing a resolution, and disseminating the updates once completed. Together these provide a documented system of risk minimization for the vendor, and also provide users with enough information to evaluate risks in their systems as well. Together they give a solid starting point for European regulators.

6.2 Incentivising vendors

It took the IT industry over a decade to adopt the model of a patch cycle driven by responsible vulnerability disclosure. Twenty years ago, vendors tried to to hush up details of breaches, using NDAs, legal threats and public relations techniques. This led to bad outcomes with researchers anonymously disclosing vulnerabilities publicly to pressure vendors into fixing them, followed by malicious exploits and emergency patches. By about ten years ago a consensus had emerged but it took several years more for the laggard firms to adopt a modern patching cycle. Vendors of products such as cars, airconditioners, industrial control system components and medical devices are still in the place that IT vendors inhabited at the turn of the century.

If nature is allowed to take its course then we can expect over a decade of repeated breaches and uncontrolled vulnerability disclosures, affecting safety-critical devices such as cars and medical devices, as well as critical infrastructure such as power grids. Vendors can expect frequent expensive recalls and system operators can expect that attacks will force them to make unplanned investments in countermeasures and compensating controls.

Regulators should move equipment vendors, systems integrators and operators towards standard processes for reporting, disclosure and patching.

Regulators need to think about how a scheme creates incentives within the relevant market. The usual incentive is the reduction of liability for compliant systems; alternatively one might look for ways to increase liability for firms that sell large numbers of network-attached devices without any means of patching vulnerabilities remotely. Perhaps these incentives will arise eventually anyway (if vulnerable cars have to be returned to a garage to have their software reflashed, this will cost enough that the car maker may support remote patching for future models) but regulators should really be leading such changes.

6.3 Establishing trust and confidence

Computer security is complex and scary; humans have evolved to be wary of adversarial action, especially when it is not well understood. Computer security fears impose real social costs, as cybercrime studies have shown; scare stories about malware and fraud online lead many people (particularly the elderly) to avoid banking or shopping online, which imposes real costs.

One of the significant social roles of the data protection authorities, for example, is to provide some reassurance about privacy so as to make citizens comfortable about online shopping. Similarly, safety regulators in the Internet of Things will earn their living by providing comfort and thus facilitating the uptake of new technologies.

The trust which users place in safety regulators must be well-founded. Yet legacy regulatory systems are falling badly behind. We already noted that medical device regulation often appears to protect the vendors' interests rather than the patients'. Second, in financial consumer protection, PCI-DSS is designed to protect banks, but not protect either consumers or small business from secondary victimisation – a point recently illustrated by a letter written by the US National Retail Foundation [14]. Third, public confidence in vehicle type-approval regulation has been shaken recently by the Volkswagen scandal.

To cut through this mess, the European Commission should move boldly to open, standardised vulnerability management across all sectors where security is becoming a critical component of safety. We know what security best practice is: security is dynamic, so patching must not just be done, but be seen to be done. The best practice of the IT industry is now incorporated in international standards, namely ISO 29147 and ISO 30111. Regulators should push firmly for their adoption everywhere.

Our next recommendation is therefore

Recommendation: The EU should encourage the adoption of a secure software development lifecycle, with a documented vulnerability management process. Safety regulators should be empowered and encouraged to pressure companies into respecting ISO/IEC 29147:2014 and ISO/IEC 30111:2013 at a minimum.

Once systems become dynamic, and are capable of being changed quickly by means of software upgrades, safety regulation must be dynamic too. The existing vulnerability management standards have evolved and proven their worth in the IT industry for a decade now. They are the best, and perhaps the only realistic, place to start.

6.4 Collecting and publishing data

The patching cycle we propose should be public to the greatest possible extent. An essential security signal to the market is the rate at which flaws are found, and the speed with which they are fixed. Rather than holding vulnerability information between ENISA and Member State security agencies, in a safety regime it must be made available to all, except in cases where there are compelling reasons to delay disclosure while critical systems are patched.

To improve a system, we have to be able to measure it. Vehicle safety is much better than fifty years ago not just because of sweeping legal measures such as product liability and type approval but because of detailed accident data that are used by multiple stakeholders. Insurers use accident statistics to set premiums; local authorities use them to prioritise road improvements; car vendors tear down crashed vehicles to see what they can learn so as to make future vehicles more crashworthy.

In a similar way, statistics on crime have driven successful crime reduction programmes in many countries, reducing the social costs of crimes from piracy through homicide to bank robbery, car theft and now cybercrime.

There are already significant levels of electronic crime against cars, in that electronic car theft tools which exploit the weak cryptographic and security mechanisms in motor vehicle remote key entry systems and immobilisers are used in significant numbers of thefts of expensive vehicles from the EU. We can see this as the start of ‘IoT crime’, and we can expect much more of

it. European institutions have a duty to collect decent statistics, so that all can understand what's happening and those stakeholders with the ability to make relevant changes can do so.

6.5 A caveat

We have seen a consumer electronics company pushing a mandatory update to its customers that did not fix a vulnerability in that was being actively exploited, but rather a vulnerability in its digital rights management system. The company was motivated to protect its business model, but not to protect its customers.

Some care may need to be taken that vendors do not exploit IoT regulators in this way. For example, if road safety regulators make it mandatory for cars to have up-to-date software in order to pass their annual roadworthiness test, then mandatory updates must include fixing safety and security problems above a given level of severity. They must not include business-model updates that downgrade the user experience. If a car vendor tried to insist that a flaw in the brakes will be fixed only if the customer accepts a software 'upgrade' that also causes the car radio to play ads whenever the car is stopped at a red traffic light – regardless of whether the driver wants to listen to ads or not – then the whole system will be brought into disrepute. Yet this is exactly how vendors can be expected to behave, if regulators let them get away with it.

7 Are flaws foreseeable or discoverable?

Establishing liability involves the demonstration of negligence and the foreseeability of harm. In this section we attempt to clarify how these principles apply to security and privacy. A fundamental question is the extent to which ‘zero days’ are foreseeable.

7.1 Novel attacks and accidents

Security, privacy and increasingly safety are problems of co-evolution. Criminals, security researchers and other opponents research and develop new attack techniques. In addition, users may by chance discover some combination of inputs to a system that causes it to behave in a surprising and exploitable way. Some of these vulnerabilities are merely tactical, while others open whole new fields of offensive security research. To quote Sergey Bratus:

“ ‘Zero-day’ means new, not known before. Any scientific result worth publishing is “zero-day” – previously unknown, just discovered. Science is pursuit of “zero-day” discoveries. Since we are computer scientists, our discoveries take the form of programs: “zero-day” programs.

Without “zero-day exploits”, claims of new security phenomena – new vulnerabilities, new types of vulnerabilities, new risks – remain hypothetical. The industry cannot waste its effort on hypotheticals. Even if they wanted to, how would they know which hypotheticals are actually worth their effort, which is necessarily limited?” [9]

The discovery of a new type of vulnerability is actually scientific discovery, the eureka moment accompanied by an exclamation of “Woot!” (This exclamation, used in hacker communities, is short for “Wow, loot” and goes back to the days of Dungeons and Dragons, before the Internet.)

To what extent can we consider zero-day exploits to be a foreseeable harm? We might naïvely assume that if an operating system has on average eight vulnerabilities patched every month, a court might consider the discovery of new ones to be entirely foreseeable. But what about innovative new types of attack?

7.2 Tool construction and scientific discoverability

Vulnerabilities tend to come in classes, with both their discovery and exploitation being linked to tools. Stack-overflow vulnerabilities were discovered in the 1960s and first exploited at scale in the Internet worm in the 1980s. During much of the 1990s, their discovery and exploitation were both a craft activity; researchers found them by staring at code, or by playing with it, and exploiting them involved writing fiddly software. Over time these activities became industrialised: fuzzing tools were written to find them at scale, and knowledge of how to exploit them became widespread. By the early 2000s, static-analysis tools were starting to appear that enabled vendors to fix not just one vulnerability, but to find and fix all similar holes in their code.

Software vendors also developed standards for secure coding; procedures for reviewing code before it was added to a product; and extensive suites of tests (including fuzzing). In well-run software firms these are now integrated into a daily development cycle whereby newly-written code is reviewed by another programmer, then checked by a static analysis tool such as Coverity, then integrated into a nightly build, and subjected to extensive testing overnight. If a firm does not adhere to this established best practice, and a vulnerability in their code leads to an attack that causes harm, a claimant can argue negligence by the standards of the industry.

One problem is sectoral diversity: software engineers in the medical device industry, for example, do not work to the security engineering standards enforced at Microsoft or Google or Apple. Infusion pump vendors have simply not yet had to deal with repeated capable attacks on their systems. Medical-device company lawyers would argue that they should be judged by the standards of their peers, not by comparison with Apple or Microsoft.

Second, we can divide zero-day vulnerabilities into two classes: those that can be found with automated tools because they are of a well-defined and studied type, and the sudden, innovative, eureka-like discoveries. For example, stack overflow vulnerabilities were mitigated for a decade using address space layout randomization (ASLR), until the sudden discovery of attacks based on return-oriented programming (ROP). Novel vulnerabilities might in some circumstances attract less liability for a period of time (product liability is strict, so a medical device maker whose product kills a patient because of a novel attack might still be held liable, as they designed a device whose software could kill people when they could have designed it differently – but we are aware of no relevant case law).

However, the great majority of exploited vulnerabilities are wholly foreseeable. The OWASP top 10 [39] is a list of the most frequently discovered web application vulnerabilities; all of these are entirely predictable, reproducible, and endemic in today’s web environment. Moreover, developers might choose from a number of readily created open source or proprietary tools to discover them. Another example of well-understood classes of vulnerabilities might be the CWE database, which one of the authors reads regularly to know what to look for on security assessments [36]. While one might be tempted to create another database here in the EU, we would encourage people to recognise there are many of these around the world already with capable motivated maintainers.

A better focus is to insist that if a tool that would find a vulnerability predates the release of a product containing it, the vendor should not be able to rely on the claim that it was “unforeseeable”.

This matters because vendors experience real initial costs when they start to use static analysis tools, which are not limited to the obvious costs of buying the tools and training their staff. As we noted earlier, a vendor who starts to use a static analysis tool to check its code base may suddenly find several thousand new bugs, and have to spend weeks fixing them. The same applies if a vendor decides to use only software developers who are chartered engineers, rather than outsourcing the coding. There is therefore an incentive for all but the best-run and most safety-conscious companies to just ignore the bug reports and hope for the best. The regulator has a real role here in preventing a race to the bottom.

7.3 Demonstrating harm

As already mentioned, metrics of harm are problematic in information security. Most of the stakeholders have an interest in overplaying harm; firms talk up the damage in order to get police to chase perpetrators, police buy this in order to justify budgets and get longer sentences for convicted offenders. The introduction of breach-disclosure laws in the USA has engaged insurance companies: writing to millions of customers is expensive enough to give rise to a claim; and insurance loss assessors are at last a significant actor with an incentive to argue the costs of a breach down as far as possible.

The Internet of Things may make loss assessment more straightforward, as there are well-understood methods for assessing the damages payable in the event of a car crash or a medical accident. Claims for disruption of

industrial production following interference with control systems are also relatively tractable using traditional cost accounting methods.

The real challenge may be whether anyone will put the time and effort into linking the causal events that lead up to a security incident, and preserving the chain of evidence.

7.4 Attribution, causation and liability shifting

There is often more heat than light generated by discussions of attribution and causation with respect to hacking incidents. Sometimes the perpetrator can be found; sometimes they cannot. Vulnerability researchers find fault with manufacturers, manufacturers blame negligent users, and users may respond by blaming the hackers. Everyone tries to shift the liability.

The temptation for regulators who are confused by and anxious about technology is to call for more user education, rather than tackling vendors or system operators. One of the roles of better regulation will be to analyse agency and causation in complicated cases not just legally but conceptually, philosophically and scientifically.

Air traffic investigations regularly accept multi-causal reports about adverse safety events, and hacking is no different. Rail, maritime, and nuclear accidents also follow this pattern; grown-up sectoral regulators should in theory be able to add the risk of malicious action, exploiting hazards caused by careless design, to the set of hazards and risks they deal with.

Let us illustrate with a simple case. An attacker send a ‘spear-phish’ to an employee of a power company, compromising their laptop. From here the attacker chains their way in, compromising the corporate mail server, using this to take over the PC of a system administrator, and finally getting the password needed to log on to an operational system in another network and in a country halfway around the world, allowing him to switch off a hydro-power plant.

If the attacker can be found and prosecuted, or if the country for which he works can have blame attributed in the world’s media, then agency is assigned. However, this story is incomplete. Some proportion of blame can also be applied to the power company for not training its staff properly and filtering incoming mail adequately; perhaps to the vendors of the laptop, PC and mail server, particularly for vulnerabilities of known types, or once more to the power company if these vulnerabilities had been patched but the company had not applied the patches; and to the designers and operators

of the hydro power plant for building and operating it in such a way that it could be closed down by the remote action of an untrusted party.

Attackers are often not caught, and so the regulator must place some weight on an effective patching cycle, which means not just coordinated disclosure and regular patch production, but also the timely application of patches by critical users. An understanding of criticality is important. As our example shows, some individuals are also critical, in that their compromise (or the compromise of their computers) can do major harm. The details of deployment also matter: some configurations are so risky that they should not be used in certain contexts. If you install typical industrial control system devices without firewalling them from the Internet, you are almost certainly negligent.

Context matters too: it is unreasonable to expect pacemaker patients to firewall their implants, so such devices must not be openly accessible. Reading a pacemaker's data and changing its configuration should only be possible once it has been strongly authenticated to a trustworthy controller.

Lastly, a naïve user might expect someone else to do the work. We saw above that medical device vendors assumed that hospital networks would be secure, so their devices would not have to be. The vendors of industrial control system components assumed that factories, refineries and other industrial plant would have entirely private networks, firewalled off effectively from the Internet. In other circumstances people have hoped that their employer, their government, or even someone else's government would block bad traffic.

Regulators must vigorously challenge such wishful thinking, before attackers do. In many cases, network attacks are deliberately indistinguishable from legitimate traffic, and it would be unreasonable to expect ISPs to block them. Even in cases where we can find some fault with the network, such as when Charlie Miller and Chris Valasek discovered that the GSM network allowed them access to vehicles on the road, the vulnerabilities are often clearer in hindsight.

Just as weak managers bully their staff and suck up to their bosses while strong ones defend their staff and stand up to unreasonable demands, so also it is basic security economics that firms with no market power will hope that someone else will protect them upstream, but do nothing to protect those downstream from them in the supply chain unless they are compelled to do so. (A firm with market power may behave 'better', but market power brings its own problems.) Regulators cannot be everywhere, and transactions involving

intermediate goods are generally left to the market. However regulators must insist on safety, and the security that is necessary for that – even if this results in regulatory action propagating up the supply chain, for example from car vendors to the manufacturers of engine management units, and even of the vendors of the software platforms they use.

This leads us to the proposition that cybersecurity has many of the aspects of a public good. The safety and security of widely-deployed and safety-critical objects such as cars and medical devices has a strong claim to be such, as a systemic compromise could widespread damage, injury and death; it is thus a public good in the sense that national defence is, and should be treated with comparable seriousness. This includes not just the soft of missions undertaken by ENISA and by national SIAs, but the regulatory and market conditions for the software on which we rely to have a secure development lifecycle and be capable of remediation when serious safety or security flaws are found. The public-good doctrine is gaining ground in security regulation and strategy [40, 7, 43].

This raises practical questions of who does the regulation and whether the structure of these institutions needs to change.

8 Cybersecurity as a public good

Many goods sold in the EU must carry a CE mark, which amounts to a declaration by the manufacturer that it complies with applicable Directives and regulations.

8.1 Who are the mainstream regulators?

Mainstream regulators in Europe range from completely centralised, to completely devolved to Member States' national authorities. Furthermore, the system is dynamic, with areas of recent innovation being less harmonised.

Motor vehicles fall under EU Whole Vehicle Type Approval (WVTA) which covers the entire vehicle (some of its constituent parts may also carry a CE mark). Automated vehicles are more complicated, and do not have approval across all Member States. The UN's Vienna Convention on Road Traffic (which sets the rules in 73 countries) was amended in 2014 to allow automated steering at speeds of up to 10 km/h, but the UK allows 25km/h (autonomous vehicles can drive anywhere in the UK, which is not signed up to the Convention). The Dutch and the Germans are keen to move driverless cars forward, but other member states don't want to have such vehicles until their safety is better proven.

Insurance and policing will be affected too. What happens to the insurance industry once accident claims are no longer based on negligence, but on product liability, which is strict? Will insurance be bundled along with the sale or lease of a car? What about collisions between driverless and driven cars? How does one stop a driverless car to interrogate a passenger? Traffic risks (pollution) are different than accident risk, and platooning (where groups of cars act in concert) might reduce accident frequency and emissions, but kill more people when accidents do occur.

Broader classes of transport regulation within the EU is even more fractious. The EC's web page lists 26 different regulators, and many more associated bodies for rail, air, maritime, and goods. Coordinating them is clearly hard enough already; what will be involved in making them aware of security issues? Is it feasible to expect each of them to hire even one security engineer to inform their policymaking and assist with standards and certification?

Medical devices are covered by three directives (status in September 2016):

- Council Directive 90/385/EEC on Active Implantable Medical Devices (AIMDD) (1990)
- Council Directive 93/42/EEC on Medical Devices (MDD) (1993)
- Council Directive 98/79/EC on In Vitro Diagnostic Medical Devices (IVDMD) (1998)

Multiple national regulators still operate, and also collaborate with non-EU regulators such as the FDA (whose approval decisions may be taken more or less as authoritative by the UK's Medicines & Healthcare products Regulatory Agency (MHRA), for example). As noted above, medical device regulators are relatively toothless, although during recent revisions, the ability for Notifying Bodies in individual Member States to make “unannounced inspections on the factory of device producers” was established. This could create an opportunity for security engineers to ask difficult questions about the secure development lifecycle and quality assurance practices of the manufacturer of in vitro medical devices.

Energy regulators meet regularly as the Council of European Energy Regulators (CEER). Not only does the group use this as a method for interfacing with the EU directly, but also for interfacing as the EU Energy regulators internationally. It has a sister organisation; the Agency for the Cooperation of Energy Regulators (ACER), which has both permanent staff and staff seconded from national regulators. This might be one place to embed security engineers to work on protecting users from outages and frauds of various kinds (a very different goal than the national defense mission of ENISA).

One serious deficit in standardisation has become notable in the smart meter projects adopted by most Member States following the Electricity Directive. Standardisation was left to Member States, and some did better than others; For example, eight of the nine Spanish distribution network operators set up the Prime consortium which adopted meter standards that led to meters costing only Eur 40 per unit. The UK project has been stalled for years in arguments over standards between nineteen committees. As a result, we lack effective European standards for the home area network – the segment that connects the smart meter to washing machines, dryers and water heaters. As a result, there is not a large enough market for the smart

appliances that would make the ‘smart home’ a reality. Such a standardisation effort would have had to be led by a European agency, in collaboration with vendors and with research support; it would have needed expertise in networking, security, usability and energy management.

8.2 Who investigates safety and security incidents?

Imagine that a terrorist takes control of a driverless vehicle, and runs it through a street filled with pedestrians. Such incidents have happened in vehicles with drivers, but police investigations of such cases are straightforward. Who will provide the expertise to support the police once autonomous vehicles start being used for serious crime and terrorism?

What other regulations will be called for? The police will want a way to disable rogue vehicles remotely, and insurers will need a method to forensically investigate what went wrong. Parents will want to give children limited use of vehicles: OK to be fetched from school, but not OK to go into town on a Saturday night unless a parent gives express permission. If a child figures out how to hack the system and get the car to go places that daddy didn’t foresee, that’s also a security breach that will need to be reported and investigated.

The point is that the definition of ‘security’ for everyday things is not at all straightforward, and will evolve over time in response to incidents and accidents.

Our second point is that there are going to be a lot of incidents and accidents. The rapidly growing use of communications media such as Facebook, WhatsApp and even old-fashioned email has led to such media being involved in very large numbers of criminal cases, as well as in less serious incidents such as cyberbullying at school. This in turn has created large volumes of demand for access to data. This in turn has led the big service companies to set up large help centres; caused law enforcement to train many officers to act as contact points for getting data; and led to legislation in many countries to overhaul laws governing law-enforcement access to data.

The big problem here is jurisdiction; it’s annoying if even a small-town police officer needs to go through a cumbersome Mutual Legal Assistance Treaty (MLAT) process to get data to pursue a cyberbullying case. As a result, some countries are passing data localization laws, requiring service firms to keep data likely to be of use to law enforcement within the jurisdiction (or easily available from it). Will we see this in the IoT as well? Will

an autonomous vehicle have to keep relevant logs on servers within the EU to pass safety and security certification?

Investigating the Internet of Things will need to become as commonplace as investigating people. This leads to the following questions:

- Who will investigate adverse events and security incidents?
- Which state bodies (police, safety regulators, security regulators) need to be capable of which safety, security, and privacy investigations?
- How do we prevent or at least manage the resulting turf wars, such as the ‘equities issue’ between SIAs who want access to everything and safety regulators who must remain transparent?

Security is about power; it’s about determining who can do what to whom. Increasingly, computer and communications security will be woven into the fabric of everyday life. The software in that fabric is not just a tool for efficiency and innovation; it is also how we control the scaling of societal harm.

Security researchers must continue to protect the public good and help consumers to defend their rights by understanding what’s going on and bringing it to public attention. Regulators must build on this and never lose sight of their mission. As they interact with big business, SIAs and police in the information security ecosystem, they must remain aware that each pursues different goals and optimises different objective functions. To make sense of this, an approach based on security economics is the natural one for a regulator to take.

8.3 Software patches as optimised product recall

One role of regulators is to realign incentives by forcing manufacturers to recall and fix dangerous products. Since the NHTSA started ordering vehicle recalls a generation ago, this has been the big stick that forced vendors to pay due regard to quality and safety. A recall of millions of vehicles caused real pain to shareholders in the way that an individual customer seeking a repair or replacement under product-liability law does not.

Part of the promise of software is that it can be upgraded remotely. Rather than writing to a million customers inviting them to bring their car in for repair, a repair can be downloaded over the air and installed next time

the engine is turned off. Provided regulators remain firm in requiring hazards to be fixed, manufacturers will already have an incentive to set up the machinery to support a patch cycle. Regulators should seek to maximise the benefit. There are some complexities, of course, and regulators should seek to understand the opportunities and constraints of shipping patches and mitigations at scale. But the benefits are so large and the evidence for this from products such as PCs and phones is so overwhelming that regulators should start to insist that all smart devices with network connectivity be capable of having safety and security flaws fixed remotely by software upgrade.

8.4 Meaningful certification for cross-border trust

Different sectors will wish to certify safety and security differently. Some will certify the software, others the teams that manage it. Some will certify the engineers who write the code and others will certify the organisations they work for. All of this has been seen in the history of certification. The goal of any EU effort though, should be to increase the confidence of the internal market, and in particular trust in the whole ecosystem of the Internet of Things.

This does not mean a doctrine of prevention and rigid perfection. Rather it means a doctrine of transparency and of continuous improvement. Customers should know what they're buying and what the risks are, as they do now to a reasonable extent with cars and medical devices. The incentives in the market should lead to things getting better and safer all the time by dealing with hazards and threats as they arise.

8.5 Maintenance and evolution of whole systems

If we see cybersecurity as a public good the next question is, which actors are responsible for proving it? Who is responsible for IoT devices when they are bad actors in a network? In other words, who is the security maintainer?

Historically the owner of a device was responsible for maintaining it; you sharpened your own sword and fixed your own wagon. As time went on and technology became more complex, vendor after-sales organisations and third-party maintainers have started to play a role, along with regulators. You have to have your car inspected every year, and it's now so complex (and electronic) that you need a garage with the appropriate equipment. Regulators have to work to ensure that manufacturers don't lock customers

in to their own dealer networks but make the necessary maintenance software available to their competitors.

Meanwhile, the trend with electronic equipment is that there are fewer and fewer user-serviceable parts inside. A modern mobile phone can have a broken screen replaced, but that's about it.

Some IoT devices are deployed cheaply and disposably, but the emergent system they make up may not be so disposable. It must not be simultaneously insecure and immortal, or we will simply have socialised the risk. One of the authors asked whether governments might let contracts centrally for specialist firms to clean up malware [12], and such policy questions will also arise in the context of IoT.

As time goes by, patching alone may not be enough. In a world of complex systems, we can expect more incidents where (as with infusion pumps) each vendor can blame others for a safety incompatibility that kills. It will not be sufficient to certify the safety and security of individual components; we have to test, certify and monitor whole systems. It is already accepted that we certify a whole car, not just its component engine, brakes, steering and so on. It is also accepted that driver training and road design are linked standards. Similarly, once we have millions of autonomous, semi-autonomous and manually-driven vehicles sharing the roads, the safety authorities had better have the authority to look at the whole picture.

Recommendation: European regulators should move from certifying products to the assurance of whole systems including how the products are used and maintained, and how the overall system evolves.

9 Cyber-covigilance: a new proposal for a multistakeholder approach in Europe

This paper has described how Europe has a number of functional regulators and standards bodies, notably ENISA (for Internet security), the European Data Protection Supervisor (for privacy) and the European standards organisations. It also has a multitude of sectoral regulators and standards agencies, of which we have described some of those active in road transport, medical devices and energy; there are more in other sectors.

All these regulators and standards bodies will be challenged by the move to embed computers and communications in everything, a move currently known as the ‘Internet of Things’ but which has been gathering momentum for a number of years. The combination of computers and communications opens the possibility of cyber-attacks and raises all the issues already familiar from attacks on PCs, phones and the Internet infrastructure.

Safety regulators now have to consider security too, and many are struggling for lack of expertise. Worse, the industries themselves are struggling; until a few years ago, no car company executive had any idea that their firm needed to hire a cryptographer. The same goes for makers of medical devices and electrotechnical equipment. So the regulators should not expect to access this expertise from among the engineers that manufacturers send to sit on their technical working groups; they simply don’t have the expertise that they ought to have, and in many cases are not even aware of the deficit.

9.1 A European Security Engineering Agency

We propose that the EU create a European Safety and Security Engineering Agency to provide a shared resource for policymakers and regulators, or expand the role of an existing agency to fill the gap.

Its mission will be to:

- support the European Commission’s policy work where technical security or cryptography issues are relevant
- support sectoral regulators among the European institutions and at the Member State level
- develop cross-sectoral policy and standards, for example arising from system integration

- act as a clearing house for data gathered via post-market surveillance and academic studies
- work to promote best practice and harmonisation
 - work to harmonise standards of people, protocols, devices, definitions, sectors, and organisations
 - work to harmonise methods of product safety, security, and privacy testing
 - work to harmonise vulnerability databases worldwide, and de-duplicate/disambiguate them
- act as a counterweight to the national-security orientation of Member State security authorities

Our principal recommendation is therefore:

Principal Recommendation: The EU should create a European Safety and Security Engineering Agency with high-level technical expertise in security engineering, safety engineering and usability engineering, that collects reports on breaches and vulnerabilities, and supports both sectoral regulators and policymakers. Alternatively, it could expand an existing agency to fill the gap.

9.2 Post-market surveillance

One of the reasons we need a European Security Engineering Agency is for collecting breach and vulnerability information that falls outside the current ENISA / NIS mechanism, whose focus is on collecting data relevant to national security. We need to collect data relevant to fraud, consumer protection, competition, environmental protection and a number of other policy goals. In essence, such investigations need to become ‘data driven’, with the focus placed where the most harm is occurring.

At present, there is almost nothing. The Banque de France runs an Observatoire on payment fraud in the Eurozone, but that is essentially a private initiative, with data not being made available beyond central banks. In general, post-market surveillance should include users and defenders of systems, networks, and devices. It should also make its data available, under NDA if need be, to bona fide researchers, so that research results are repeatable.

9.3 Use of existing labs for post-market studies

Some industries have existing accredited facilities for testing compliance with (e.g.) motor vehicle type approval. Can we draw them in to become part of the future cybersecurity system?

We believe they have an important role to play in post-market surveillance: as safety and security become intertwined, we need safety labs to learn more about security, and vice versa. Digital forensics labs can report on commonly used exploits and vulnerabilities; penetration testers can share experience of common methods and techniques; safety engineers can highlight adverse events that are unexplained or may have had malicious manipulation; and the two tribes, of safety engineers and security engineers, can gradually be brought together. The best way to do this is to get them working together on shared problems. This may be the only way to effectively embed adversarial thinking into the community of safety engineers and test engineers who advise the existing sectoral standards and certification bodies.

9.4 Proportionality in causation of harm during investigations and legal remediation

It is not reasonable to apportion all blame to a single stakeholder during software security, privacy, and safety investigations in every case. Therefore, principles of proportionality of negligence need to be established, and put into practice across the EC. This will involve lengthy discussions with legal teams, and technologists, but should be started early. Eventually, Member State and Sectoral Regulators will each individually adopt this approach with respect to their constituents.

9.5 Summary of recommendations

We list here on one page our recommendations.

Principal Recommendation: The EU should create a European Safety and Security Engineering Agency with high-level technical expertise in security engineering, safety engineering and usability engineering, that collects reports on breaches and vulnerabilities, and supports both sectoral regulators and policymakers. Alternatively, it could expand an existing agency to fill the gap.

Recommendation: The EU should update the Product Liability Directive to cope with systems that incorporate multiple products and services.

Recommendation: Vendors of devices containing software should be required to self-certify, as part of the CE process, that their products are secure by default, and where these devices can be attached to a network they should by default be capable of online software update so that safety and security flaws can be fixed.

Recommendation: The EU should update the Network and Information Security Directive so that breaches and vulnerabilities are reported to affected users and to relevant regulators.

Recommendation: The Commission should investigate how to move European standards bodies in the direction of considering safety and security together in future standards work.

Recommendation: The EU should encourage the adoption of a secure software development lifecycle, with a documented vulnerability management process. Safety regulators should be empowered and encouraged to pressure companies into respecting ISO/IEC 29147:2014 and ISO/IEC 30111:2013 at a minimum.

Recommendation: European regulators should move from certifying products to the assurance of whole systems including how the products are used and maintained, and how the overall system evolves.

10 Illustrative energy sector case study

The Dragonfly campaign of malware showed up throughout Europe (and the US) throughout 2013-2014, and affected multiple European countries including Spain, France, Italy, Germany, Greece, Poland, and Romania [44]. It was given different names by different security vendors: Energetic Bear [29] and Havex [22] (which illustrates our point on standardising naming and incident reports).

We will first describe this campaign of attacks on the energy, oil and gas, and manufacturing sectors. Then we will consider how our recommendations might have mitigated it, had they been already in force.

10.1 Real case study

The attackers ran a spam / spear-phishing campaign from February to June 2013, though presumably their reconnaissance of target companies began earlier. The most popular exploit used was a Flash vulnerability already known at the time, CVE-2011-0611. However there were four others used too, including two that were zero-day (unknown) at the time: Java 6 & 7 exploits CVE-2012-1723 and CVE-2013-2465, and Internet explorer 7 and 8 vulnerabilities CVE-2012-4792 and CVE-2013-1347. Although these were exploits of commodity system software, they led to exploits of industrial systems. This is a common pattern.

The malware also made extensive use of Remote Access Trojans or Remote Administration Tools (RATs) that were implanted by this spam campaign. According to F-Secure, there were 88 variants of the RAT [22], communicating with any of a number of 146 command and control (C&C) servers that were used to control the infected machines.

Later in the campaign, the attackers changed tactics, using redirection attacks against websites, which among other things compromised the update mechanism of ICS vendor software. These vendors were then used to compromise and infect their own customers, who innocently believed they were downloading trustworthy updates.

Once inside the victim companies, the attackers began scanning for OPC clients and servers. OPC itself is a standard communications protocol for sharing real time data from industrial facilities [52]. Thus the attackers were interested either in the data, or in the servers as an entry point into the industrial networks themselves.

The multiple reports cited imply that the campaign affected the low hundreds of companies; the exact number is unclear. The different vendors report different numbers, and it is not possible to disambiguate them to get an accurate lower bound on infections. It is entirely possible other infections went unnoticed entirely; victim counts are almost always a lower bound. This holds for government analysts as well as academic researchers.

Let us examine how our recommendations might have made a difference.

10.2 Counterfactual cyber-covigilance case study

Assume that a European Safety and Security Engineering Agency (ESSEA) had already been established in 2010. They work with local member state regulators, and also offer a number of channels for existing labs and security service companies to report vulnerabilities. These surveillance reports come from a variety of partners, both altruistic and profit-driven. Safety testing labs note adverse events with devices under test, some of which lead to security investigations. Private penetration testing companies share frequency counts of CVE and CWE discoveries from tests in return for early access to aggregate data. With this as a background, let us replay the events of 2013–2014 in this alternate universe.

During 2013, aggregate data from multiple sources show a slight rise in the use of Flash vulnerability CVE-2011-0611. Because ESSEA created robust naming conventions, reporting procedures and analytics, they have helped harmonise industry reports and share feeds with US service firms, US state sector bodies such as the NCFTA and government labs in Brazil, India and Japan. Security professionals can share information, knowing they are talking about the same vulnerabilities and the same attacks.

The slight rise in CVE-2011-0611 exploitation is noted to include spear phishing emails to energy companies, and so ESSEA contacts energy regulators to investigate further. Since Adobe has already issued a patch for the Flash vulnerability, liability has shifted to their users. But Adobe continue to help investigations and try to drive wider adoption of their patch.

When the Java and Internet Explorer exploits are found on multiple websites that have been taken over by the attackers, ESSEA work rapidly with computer emergency response teams in countries hosting the sites to get them cleaned up. Oracle and Microsoft produce patches promptly, as ESSEA engineers are on first-name terms with abuse team leaders at the major vendors. Additionally, law enforcement agencies in the countries hosting the command

and control servers are informed, and victim identification operations begin in earnest.

The ICS vendors with compromised websites and downloads have been pushed by their regulators to adopt vulnerability reporting in line with ISO/IEC 29147 and now have a proper process to fix their download mechanisms, notify their customers, ship patches for their controllers and finally go public with news of the breach. In parallel they support an incident investigation by ESSEA engineers and national authorities. Victim notification goes smoothly, and the forensic evidence is preserved for analysis by Europol, ESSEA and national police forces.

ESSEA engineers were aware that OPC software companies also needed to perform coordinated vulnerability disclosure policies, and so these firms were notified. Under the current system they would not have been considered ‘critical infrastructure’ and thus would have been outside ENISA’s remit.

It is in such complex but realistic events that ESSEA would really show its value. Its engineers would be able to tell proper security engineering from wishful thinking, and to push industrial systems firms to use the appropriate standards for cryptographic signing of firmware updates rather than just something that one of their programmers thought up one Friday afternoon.

Over time, the different firms in the supply chain, from the ICS vendors to their OPC software suppliers, converge as the industry moves towards a new software security lifecycle to avoid liability. Since the OPC issues that were exploited could have been discovered by a static analysis tool, they were very clearly foreseeable. Those OPC vendors that can assure their ICS customers that they use a proper development cycle thrive, while the others – those who skimmed on the cost of hiring proper engineers and using static analysis tools – go to the wall.

The EU and Member States get a clear picture of the cost of the incidents, both to individual businesses and to the European economy. This analysis leads to better policy, both on defence and on safety, and a better understanding of how to manage the tensions between them.

Insurance companies compensate a few businesses who were harmed, but their pricing remains reasonable, because they have now understand the risks better and are confident that both risk correlation and uncertainty are being steadily reduced. This is because they get ever more data about the cost of incidents resulting from each single vulnerability over time. This knowledge in turn informs security research into what kind of coding errors are the most costly overall, which in turn leads to better static analysis tools.

Although some of the data are kept confidential while the patching is in progress, within a few months ESSEA and other involved agencies are in a position to publish full technical details of the attack along with aggregate data about its impact. This forms an open annex to a final incident report that can be read by all, so that all can learn from it.

Such transparency is the only way we can manage a world of increasingly complex attacks, where several combinations of vulnerabilities are used on hundreds of victims in a coordinated campaign. Not all the victims suffered serious harm, but we need to learn from all of them if we are to forestall similar harm to others in the future. Such information cannot usefully be held merely in security and intelligence agencies; developers cannot imagine all uses for their software, and companies cannot foresee how the systems they operate might be abused. The developers of Flash probably did not foresee their code being the infection point for oil and gas facilities; indeed, Adobe explicitly told developers not to use it in critical applications. It is useful that European law limits the effect of such disclaimers. This provides a needed incentive for developers to pay attention and learn from adversaries.

Just as Europe is now considered by people in Silicon Valley to be the world's privacy regulator, so also Europe becomes the world's safety regulator. The CE mark that signals goods made in Europe, or at least for Europe and to European standards, becomes a universal quality seal. Europe's brand is enhanced, to the benefit of European industry.

The key to all this, in the new world of the Internet of Things, is unifying the management of safety and security. Our proposed cyber-covigilance system accepts that future risk will emerge in ways we cannot anticipate. We call it 'covigilance', as we all watch; thanks to transparency, we can watch, and thanks to liability rules, we have an incentive to. Wherever code we have written is being used in places and in ways that could harm to our fellow humans, that is our business, and we need to know. Only that way can we all learn and improve.

11 Conclusion

As computers and communications become embedded everywhere, software will play an ever greater role in our lives. From autonomous cars to smart meters, and from embedded medical devices to intelligent cities, one environment after another will become software driven, and will start to behave in many ways like the software industry. There will be the good, the bad, and the ugly. The good will include not just greater economic efficiency but the ability to innovate rapidly on top of widely deployed platforms by writing apps that build on them. The bad will range from safety hazards caused by software bugs to monopolies that emerge as some of the apps become dominant. The ugly includes attacks. Vendors, regulators and society as a whole have to start thinking about malicious adversaries as well as about random chance; about deception as well as about unforced errors.

There will be a huge net benefit to society, but the bad and the ugly aspects will cause more work for regulators. What's more, this work will be more complex and will require new approaches.

At present, the regulation of safety in the EU consists of a few overarching laws (such as those on product liability and data protection) plus a rich and evolving fabric of detailed rules governing particular sectors. In the future, safety will require security as well. Again, there will be a few overarching laws (which must include frameworks for vulnerability disclosure and software update); plus competent and detailed sectoral regulation.

At present, the regulation of safety is largely static, consisting of pre-market testing according to standards that change slowly if at all. Product recalls are rare, and feedback from post-market surveillance is slow, with a time constant of several years. In the future, safety with security will be much more dynamic; vendors of safety-critical devices will patch their systems once a month, just as phone and computer vendors do now.

This will require major changes to safety regulation and certification, made more complex by multiple regulatory goals. A smart meter, for example, has a national-security aspect (a hostile actor must not be able to switch off millions of households at once), a competition aspect (the power company should not be able to abuse security or surveillance mechanisms to lock in customers), a revenue enforcement / policing aspect (the customer should not be able to steal electricity), and a variety of other aspects such as fire safety. For these reasons, a multi-stakeholder approach involving co-vigilance by multiple actors is inevitable.

One critical missing resource is expertise on cybersecurity, and particularly for the European regulators and other institutions that will have to adapt to this new world. An agency with the relevant expertise and remit is now essential. The EU needs to consider whether to create a new agency, or extend an existing one, in order to fill the gap.

Éireann Leverett, Richard Clayton, Ross Anderson
Cambridge and London, September 21st, 2017

References

- [1] Ross Anderson, Rainer Böhme, Richard Clayton, Tyler Moore: *Security Economics and the Internal Market*. Study commissioned by ENISA, 2008
- [2] Ross Anderson, Shailendra Fuloria: *Security Economics and Critical National Infrastructure*. In: *Economics of Information Security and Privacy*. Springer US, 2010
- [3] Ross Anderson, Shailendra Fuloria, Éireann Leverett: *Who Controls the Off Switch?* Computer Laboratory, University of Cambridge, 2010
- [4] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel van Eeten, Michael Levi, Tyler Moore, Stefan Savage: *Measuring the cost of cybercrime*. In: *The Economics of Security and Privacy*, Springer, 2015
- [5] Richard J Arnould, Henry Grabowski: *Auto safety regulation: An analysis of market failure*. The Bell Journal of Economics, 1981
- [6] P Aspden, J Wolcott, J L Bootman, L R Cronenwett: *Preventing Medication Errors*. National Academies Press, 2007
- [7] Johannes M Bauer, Michel J G Van Eeten: *Cybersecurity: Stakeholder incentives, externalities, and policy options*. Telecommunications Policy, 33(10), 2009
- [8] Matt Bishop: *What is computer security?* IEEE Security & Privacy, 1, 2003
- [9] Sergey Bratus: *Open Letter: Re: Draft report on Human rights and technology: the impact of intrusion and surveillance systems on human rights in third countries, item #17, regulating zero-day exploits*. <http://www.cs.dartmouth.edu/~sergey/wassenaar/letter-to-mep-marietje-schaake-re-exploit-regulation.pdf>, March 14, 2015
- [10] Michael G Bridge: *Benjamin's sale of goods*. Sweet & Maxwell, 2012

- [11] Pilita Clark and Sam Jones: *GCHQ intervenes to secure smart meters against hackers*. (accessed August 19, 2016) Financial Times, <http://www.ft.com/cms/s/0/ca2d7684-ed15-11e5-bb79-2303682345c8.html> 2016
- [12] Richard Clayton: *Might governments clean-up malware?* Communication and Strategies, 2011
- [13] Anupam Datta, Ante Derek, John C Mitchell, Arnab Roy: *Protocol composition logic (PCL)*. Electronic Notes in Theoretical Computer Science, 2007
- [14] Mallory B. Duncan: *Letter from the NRF to the FTC on PCI*. <https://nrf.com/sites/default/files/PCI-2016-NRF%20White%20Paper%20on%20PCI%20DSS.pdf>, May 23, 2016
- [15] ENISA: *Good Practices for an EU ICS Testing Coordination Capability*. 2013
- [16] ENISA: *Methodologies for the identification of Critical Information Infrastructure assets and services*. 2015
- [17] EPCIP: *Communication from the Commission on a European Programme for Critical Infrastructure Protection*. (accessed September 2, 2016) <http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52006DC0786&from=EN>, 2006
- [18] Scott Erven, Mark Collao: *Medical Devices: Pwnage and Honey-pots*. (accessed August 16, 2016) https://www.youtube.com/watch?v=ZusL2BY6_XU, 2015
- [19] European Commission: *Revisions of Medical Device Directives*. (accessed August 19, 2016) http://ec.europa.eu/growth/sectors/medical-devices/regulatory-framework/revision_en, 2016
- [20] European Commission: *The 'Blue Guide' on the implementation of EU product rules 2016*. 2016
- [21] Euro NCAP: *EURO NCAP – Members and Test Facilities*. (accessed September 8, 2016) <http://www.euroncap.com/en/about-euro-ncap/members-and-test-facilities/>, 2016

- [22] F-Secure: *Havex Hunts For ICS/SCADA Systems*. (accessed September 8, 2016) <https://www.f-secure.com/weblog/archives/00002718.html>, 2014
- [23] FDA: *MAUDE – Manufacturer and User Facility Device Experience*. (accessed August 18, 2016) <http://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfMAUDE/search.CFM>, 2016
- [24] Daniel Halperin, Thomas S Heydt-Benjamin, Benjamin Ransford, Shane S Clark, Benessa Defend, Will Morgan, Kevin Fu, Tadayoshi Kohno, William H Maisel: *Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses*. IEEE Symposium on Security and Privacy, 2008
- [25] Y Y Han, J A Carcillo, S T Venkataraman, R S Clark, R S Watson, T C Nguyen, H Bayir, R A Orr: *Unexpected increased mortality after implementation of a commercially sold computerized physician order entry system*. Pediatrics, 116, 2005
- [26] Erik Hollnagel, David D Woods, Nancy Leveson: *Resilience engineering: concepts and precepts*. Ashgate Publishing Ltd, 2007
- [27] Institute of Medicine: *Health IT and Patient Safety*. National Academies Press, 2011
- [28] International Organization for Standardization: *ISO/IEC 29147:2014 Information technology – Security techniques – Vulnerability disclosure*. (accessed July 11, 2016) ISO, Geneva, Switzerland, http://www.iso.org/iso/catalogue_detail.htm?csnumber=45170
- [29] Kaspersky: *Energetic Bear – Crouching Yeti*. (accessed September 8, 2016) <https://securelist.com/files/2014/07/EB-YetiJuly2014-Public.pdf>, 2014
- [30] Page Keeton: *Products Liability. Proof of the manufacturer’s negligence*. Virginia Law Review, 1963
- [31] Daniel B Kramer, Yongtian T Tan, Chiaki Sato, and Aaron S Kesselheim: *Postmarket surveillance of medical devices: a comparison of strategies in the US, EU, Japan, and China*. PLoS Med, 10(9), 2013

- [32] Nancy G Leveson, Clark S Turner: *An investigation of the Therac-25 accidents*. Computer 26(7), 1993
- [33] Paolo Masci, Rimvydas Rukšenas, Patrick Oladimeji, Abigail Cauchi, Andy Gimblett, Yunqiu Li, Paul Curzon, Harold Thimbleby: *The benefits of formalising design guidelines: A case study on the predictability of drug infusion pumps*. Innovations in Systems and Software Engineering, 11(2), 2015
- [34] Jerry L Mashaw, David L Harfst: *The struggle for auto safety*. Harvard University Press Cambridge, MA, 1990
- [35] MHRA: *MHRA Process Licensing Portal*. (accessed August 18, 2016) <https://pclportal.mhra.gov.uk/>, 2016
- [36] MITRE: *Common Weakness Enumeration*, (accessed September 3, 2016) <https://cwe.mitre.org/>, 2007
- [37] Marie Moe, *Go Ahead, Hackers. Break My Heart*. (accessed August 16, 2016) <https://www.wired.com/2016/03/go-ahead-hackers-break-heart/> 2016
- [38] Steven J Murdoch, Mike Bond, Ross Anderson: *How Certification Systems Fail: Lessons from the Ware Report*. Computer Laboratory, University of Cambridge 2012
- [39] OWASP: *Open Web Application Security Project*. (accessed September 3, 2016) https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project, 2010
- [40] Benjamin Powell: *Is Cyberspace a Public Good – Evidence from the Financial Services Industry*. Journal Economics and Policy, 2005
- [41] Monte Reel and Jordan Robertson: *It's Way Too Easy to Hack the Hospital*. Bloomberg Businessweek, Nov 2015
- [42] Karen Sandler, Lysandra Ohrstrom, Laura Moy, Robert McVay: *Killed by code: Software transparency in implantable medical devices*. Software Freedom Law Center, 2010

- [43] F Schneider, E Sedenberg, D Mulligan: *Public Cybersecurity and Rationalizing Information Sharing*. <https://www.irgc.org/wp-content/uploads/2016/04/IRGC-Public-Cybersecurity-0P-2016.pdf>, Opinion piece for the International Risk Governance Center (IRGC), Lausanne: IRGC, 2016
- [44] Symantec: *Dragon fly: Cyberespionage Attacks Against Energy Suppliers*. (accessed September 8, 2016) https://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/Dragonfly_Threat_Against_Western_Energy_Suppliers.pdf, 2014
- [45] Harold Thimbleby: *Improving safety in medical devices and systems*. IEEE International Conference on Healthcare Informatics (ICHI), IEEE, 2013
- [46] Ken Thompson: *Reflections on trusting trust*. Communications of the ACM, 27(8), 1984
- [47] Iain Thomson: *Our pacemakers are totally secure, says short-sold St Jude*. (accessed September 1, 2016) The Register, http://www.theregister.co.uk/2016/08/29/st_jude_hits_back_at_shortselling_security_firms_claims/, 2016
- [48] Dan Tofan, Theodoros Nikolakopoulos, Eleni Darra: *The cost of incidents affecting CII*s. ENISA, 2016
- [49] *US vs. Guidant LLC 2011*. 2011
- [50] Wikipedia, The Free Encyclopedia: *Barnaby Jack*. (accessed August 16, 2016) https://en.wikipedia.org/wiki/Barnaby_Jack,
- [51] Wikipedia, The Free Encyclopedia: *Ford Pinto*. (accessed September 3, 2016) https://en.wikipedia.org/wiki/Ford_Pinto
- [52] Wikipedia, The Free Encyclopedia: *Open Platform Communications*. (accessed September 8, 2016) https://en.wikipedia.org/wiki/Open_Platform_Communications
- [53] Wikipedia, The Free Encyclopedia: *Pharmacovigilance*. (accessed July 4, 2016) <https://en.wikipedia.org/wiki/Pharmacovigilance>

GETTING IN TOUCH WITH THE EU

In person

All over the European Union there are hundreds of Europe Direct information centres. You can find the address of the centre nearest you at: <http://europa.eu/contact>

On the phone or by email

Europe Direct is a service that answers your questions about the European Union. You can contact this service:

- by freephone: 00 800 6 7 8 9 10 11 (certain operators may charge for these calls),
- at the following standard number: +32 22999696, or
- by electronic mail via: <http://europa.eu/contact>

FINDING INFORMATION ABOUT THE EU

Online

Information about the European Union in all the official languages of the EU is available on the Europa website at: <http://europa.eu>

EU publications

You can download or order free and priced EU publications from EU Bookshop at: <http://bookshop.europa.eu>. Multiple copies of free publications may be obtained by contacting Europe Direct or your local information centre (see <http://europa.eu/contact>).

JRC Mission

As the science and knowledge service of the European Commission, the Joint Research Centre's mission is to support EU policies with independent evidence throughout the whole policy cycle.



EU Science Hub
ec.europa.eu/jrc



@EU_ScienceHub



EU Science Hub - Joint Research Centre



Joint Research Centre



EU Science Hub

doi:10.2760/25200

ISBN 978-92-79-77863-6



Publications Office