Standardisation and Certification of the 'Internet of Things'

Éireann Leverett[1] Richard Clayton[2] Ross Anderson[3]

May 22, 2017

Abstract

We report on a research project for the European Commission into what will happen to safety regulation once computers are embedded invisibly everywhere. The European Union already regulates many aspects of the safety of vehicles, medical devices, electrical equipment, domestic appliances and even toys. As these devices and systems are recruited to 'The Internet of Things', their vulnerabilities (whether old or new) may be remotely exploited, with consequent risks. Many regulators who previously thought only in terms of safety will have to start thinking of security as well. The systems and devices that are starting to expose their security and safety vulnerabilities to the whole Internet are certified under a disparate range of European, national, industry and other schemes. In this paper we describe the problems and outline the opportunities for governments, industry and researchers. The EU is already the world's main privacy regulator, as Washington doesn't care and nobody else is big enough to matter; it should aim to become the main safety regulator too – or risk compromising the safety mission it already has. To deliver, it will need to coordinate the 'rows' of liability, transparency and privacy principles with the 'columns' of 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, rather than having policy in these areas overshadowed or even pre-empted by Member States' national security concerns. For industry and practitioners, the main message is that safety and security are merging: safety engineers are going to have to learn all about security, and vice versa. This affects everyone from working engineers to the folks in the test labs and the regulators' committees that set the standards to which they test. Researchers will have lots of new topics, from the design of the next generation of regulatory institutions to technical topics such as sustainability of software and the toolchains that support it. How do we write code for which security patches must be made available for the next 30 years? This poses many fascinating new combinations of problems in both engineering and economics.

[1]Privacy International: eireann@privacyinternational.org
[2]Computer Laboratory, University of Cambridge: richard.clayton@cl.cam.ac.uk
[3]Computer Laboratory, University of Cambridge: ross.anderson@cl.cam.ac.uk

1 Introduction

Governments have long had an important role in maximising social welfare by regulating safety, wherever private-sector providers do not have adequate incentives to do so. The motor industry spent decades competing to decorate cars with chromium rather than fit them with seat belts, until governments took action. In the USA, Ralph Nader persuaded Congress to set up the National Highways Transportation Safety Administration and the provision of crashworthiness information; Europe followed with the Product Liability Directive and mandatory safety testing. The regulation of drugs has similarly moved from the Wild West of nineteenth-century patent medicines to modern standards of safety and efficacy assessed by randomised controlled trials. Toys can no longer have lead paint, and if you pull out a teddy bear's eye it must not leave a sharp spike behind. Regulation plays a key role in consumer confidence.

Not all regulation comes from the state; the insurance industry plays its part, with insurance premiums translating safety test results into incentives, and insurance labs producing the needed results where government labs don't. (In some cases, such as fire and burglar alarms, insurers' labs set the standards too.) Indeed, governments sometimes justify intervention in protective cybersecurity on the grounds that they are the insurer of last resort. There is also some industry self-regulation, notably through standards bodies.

In short, safety is a large and complex ecosystem of both private and public sector regulation. Like many other sectors of the economy, it is being disrupted by technology. The dependability – the safety and security – of computer and communications systems is becoming ever-more 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 many regulatory agencies concerned with safety in a variety of industries. We were commissioned in 2016 by the European Commission's Joint Research Centre to examine the effect of 'The Internet of Things' on the ecosystem of safety regulation: the core question was what the EU's regulatory framework should 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? And how do we get there from here? This paper will consider the broader lessons that we draw from our work, but we will start by considering security regulation.

The social welfare goals of a cybersecurity regulator (whether free-standing 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, law enforcement, competition, the accurate recording of transactions and dispute resolution. An example where all are in play is smart electricity meters: we do not want the meters in our homes to cause fires, to leak personal data, to enable a foreign power to threaten to turn them off, to allow the utility to exploit market power, to make electricity theft easy, or to make it impossible to resolve disputes fairly.

Achieving these goals will depend on mechanisms such as cryptography and access control; assurance that these have been correctly implemented and are being effectively maintained; standards for interoperability; and liability rules that align incentives. There must

be mechanisms to prevent actors externalising risks to third parties who are less able to cope with them. This all involves not just writing standards and testing equipment before installation, but monitoring systems for breaches and vulnerabilities, and compelling the updating of software to deal with both safety and security issues as they arise.

So the goals and mission of a cybersecurity regulator may be a mix of the following:

1. Ascertaining, agreeing, and harmonising 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.

The regulators' first task is policy: to determine what needs to be regulated, and why. There will 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 groups and relevant international bodies. Standards can build on existing specialist work, such as the US National Institute of Standards and Technology (NIST) standards for cryptography.

The first goal appears self-evident. What bad outcomes are we seeking to prevent or to mitigate? But this question can be tricky. Exactly whose risk should the regulator be reducing – the risk to a dominant industrial player, or to its millions of customers?

The second goal often entails adapting or evolving existing standards, of which there are already many, from algorithms and protocols through software processes to how people should be managed in organisations. Again, it can be tricky: inappropriate standards have been adopted in the past for both political and commercial reasons [2].

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. This is also tricky, as it creates incentives for liability games and regulatory capture.

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 that 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 will be needed to minimise the amount of time the purchaser is exposed to attacks. Online software updates also cut the costs of doing product recalls to fix safety problems. The tricky bit here is which vulnerabilities get prioritised.

The fifth goal seems to be focussed on reducing the exposure of insurers. There are other defenders too: consumers; vendors; security service companies; computer emergency response teams (CERTs); and finally, government agencies charged with protecting critical infrastructure. But some accidents are more salient than others, and voters are not content for all fatal accidents to be valued equally. If insurers or regulators act as if they are, a political backlash may follow (as happened in the USA with the Ford Pinto and in the UK with train protection systems).

The sixth goal is also about reducing asymmetry between vendors and customers. But the focus is no longer on technical vulnerabilities, but on the overall impact of externalities. If malware causes a business to lose money, the regulator might not focus on loss prevention, but ensure that the liability falls on the party most capable of mitigating the risk, such as the bank. Similarly, car companies should bear the cost of unsafe software causing accidents. If autonomous vehicles are bundled with insurance, then the incentives may be broadly in the right place, at least for the first purchaser; but the regulator may still have to consider how used vehicles will get security and safety patches, and the cost of their insurance. In fact, the EU has already created a 'right to repair' to open up after-markets for car parts; how might it adapt this to the need for security patches?

2 Historical approaches to regulation

Safety engineering is both about applications such as transport (where licensed drivers and 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 security and privacy. The security engineer's task is to enable normal users, and even vulnerable users, to enjoy reasonable protection against a capable motivated opponent. 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 and human security. It's not just about 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 abuse cases do we have to consider? And what about pensioners with Alzheimer's?

The task is to embed adversarial thinking into standards bodies, enforcement agencies, and testing facilities. To scope out the problem, we studied the history of safety and standards in various contexts: road transport, medical devices and electrotechnical equipment.

2.1 Road transport

It took much of the 20th century for road vehicle safety to be properly regulated. (Rail safety took much of the 19th.) Car manufacturers initially took the view that if you were injured in an accident you should sue the driver who injured you; and if he blamed the car he should sue the person he bought it from, and so on. 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 started from a belief that crash testing alone would be enough, but found it needed to force the recall of unsafe vehicles. The story is told in 'The Struggle for Auto Safety' [17].

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. But here, too, experience teaches that general policy measures cannot do all the work. Europe now has a substantial body of specific regulation on transport safety. Framework Directive 2007/43/EC harmonises the type approval of vehicles; many evolving safety requirements are subsumed under this Directive. Most of the technical standards come via UNECE, which includes not just EU Member States but other car-making countries such as Russia, Japan, Australia, South Africa and Tunisia; however it's EU law that gives real force to the standards, and sets the rules for quality assurance processes such as independent testing with which US makers generally also comply on a voluntary basis.

Other regulations have harmed vehicle security, most notably the Wassennaar Arrangement export controls (Regulation 428/2009) which limited cryptographic keylength, with the effect that many common remote-key-entry systems are simple to brute force – making car theft much easier than it should be.

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 – safety regulation has to take a holistic view of an application. This leads to sectoral safety regulators. In the USA, for example, the NHTSA can regulate not just the carmakers, but also the environment and the drivers (through speed limits). The European system is more distributed, with at least 26 agencies involved in car safety, but the same principle applies: it is more natural to embed security regulation in existing transport regulation rather than in a new general 'security', 'cyber' or 'data protection' law.

The move to autonomy will make safety regulation more acute; it will become more complex as the software and associated systems do; and it will become 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.

The lessons to be learned from transport are that security in this context is largely about safety; that while there are some useful over-arching measures such as those on liability and transparency, much of the regulatory work will be detailed and application-specific; and it will evolve over time as vehicles become smarter, as the environment changes, and perhaps also as people change; drivers may be less skilled thirty years from now.

2.2 Healthcare

Like motor vehicles, medical devices are safety-critical devices, with failures in their design and use being responsible for a comparable number of 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." [21]

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.

The Therac accidents in the 80s first brought medical device safety to public attention [21]. Software defects in a medical accelerator injured or killed a number of patients, and it took some time for the problem to be reproduced. These accidents also illustrate the inadequacy of relying on liability laws alone, as the US Therac victims needed to; it took five years from the initial incidents to public safety reports.

The system is still unsatisfactory. The US FDA lets devices come into circulation following only a review of documentation and without testing of their functionality and usability. The EU is also dragging its feet; medical devices are mostly regulated by Member States, but the Medical Devices Directive (2007/47/EC) requires Member States to operate vigilance systems, and the Commission is working towards harmonising the system, as they have already done for pharmaceuticals.

Recalls of defective medical devices do 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 [25]. An analysis of 23 recalls by the 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 [20].

Work by Thimbleby and others has demonstrated that the largest number of avoidable deaths involving medical devices are usability failures, and that the level of fatalities is comparable to the number of road traffic deaths [16]. 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 variants 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 other variants bearing the same model number they are 5 and 0.

Such usability conflicts are not in principle hard to tackle. Aviation regulators do not permit a pilot qualified on a Boeing 767 to fly an Airbus 340 without further training; thus 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 underlying failure 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 even include a usability study of the submitted device in isolation, let alone a more ecologically valid study of its likely use in a real hospital environment. So there is no real opportunity to come to grips with safety usability. Despite a series of complaints, the FDA appears not to have recalled a single infusion pump on usability grounds. The safety incentives are also much weaker

for medical devices than they are, say, for aircraft. While plane crashes are highly public, post-market surveillance 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 a blame game, with vendors claiming that hospital network administrators should block attacks at their firewall, while hospitals claim that the vendors should make devices hackproof. And 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 [19].

This blame game leads to muddy waters – which ironically is the name of a market research firm that shorted stocks in St Jude medical devices, just before announcing vulnerabilities in their pacemakers [23]. St Jude has filed a defamation suit against Muddy Waters, which has denounced this as an attack on free speech; the FDA says it doesn't comment on pending litigation.

A survey of implanted medical device standards and regulation noted in regard to post-market surveillance (PS):

> "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." [15]

By not permitting notified bodies [NBs] and competent authorities [CAs] to study what happens after they grant approvals, the EU has failed to collect the evidence that would be most useful to security and safety regulators and researchers alike. Of course post-approval studies must preserve privacy, but pharmacovigilance authorities already know how to deal with this. Luckily, revisions of the Medical Device Directives are under way, which may give an opportunity to improve matters [12]. We have suggested that the Commission adopt adverse event reporting along the lines of the system currently used for pharmacovigilance. This would also start to align the medical device industry with the disclosure standards normal in the IT industry, as we will discuss later.

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 infrastructure (CII): "The most affected CII sectors seem to be financial, ICT and energy" [24].

The US electric sector regulation ecosystem gives examples of what can go wrong. The USA has five interconnects, nine regional markets, roughly 1900 bulk power operating organisations across generation and transmission, and many more distribution companies. All of these companies have a role to play in securing the grid. The regulator, the North American Electric Reliability Council (NERC), required operators to have (fairly minimal) cybersecurity controls for critical assets, but did not allow them to include the expense within their regulated cost base; so operators had to pay for cybersecurity rather than passing on the bill to customers. Among the assets defined as critical was any generating plant with a 'black start' capacity: the ability to start up even if the grid is down. Fossil-fuel generators generally have black-start capability only if they have auxiliary diesel generators. It transpired that plant managers were removing these to avoid paying for NERC compliance. An attempt to make the US grid more secure against attack thus ended up making it less secure, including against random failure [2].

It is notoriously difficult to put a price on dependability; the UK is wrestling with possible designs for market mechanisms to provide surplus capacity, of which more will be needed as we move to variable energy sources. Not all of the components on which we rely for dependable energy have explicit prices associated with them.

A further lesson is in attitudes to standards. The IT industry is entrepreneurial, 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 problems when IT bodies offer multiple standards that are not only incompatible, but actually in conflict.

Another lesson can be learned from smart meter deployment. In 2009 The UK Government decided to deploy smart meters, following Directive 2009/EC/72, aiming for 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 sure that this could not be exploited by an opponent [3]. This led to a flurry of activity and GCHQ got involved in redesigning smart meter security [8]. Presentations by their officials make 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 to increase their market power or even defraud users directly. This is despite the electricity regulator (Ofgem) being concerned that fraud against customers could get worse.

The NIS Directive (2013/0027 COD) brings in quite a separate regulator. It requires Member States to arrange for firms that are part of the critical national infrastructure to report all security breaches and vulnerabilities to some central government agency. Following the smart meter case, one ought to ask whether such a body will have the mandate and the means to coordinate effectively with industry regulators.

3 Generic approaches to the problem

At present, cybersecurity regulation in Europe is a mixture of national and European. Each Member State defends against external threats using one or more security / intelligence agencies or CERTs, which are coordinated by ENISA. Privacy is the responsibility of a separate agency, the European Data Protection Supervisor (EDPS). Meanwhile the safety regulators for vehicles, health, energy etc. operate increasingly at the European level, but have little cybersecurity expertise.

Our research project's purpose was to consider where we want to be in 10–20 years' time. Do we want a single EU cybersecurity regulator 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, DG COMP to support competition, somebody else for 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 with national regulators?

We will start off by considering the generic approaches in more detail.

3.1 Liability

The software industry has fought hard for fifty years to avoid liability for its products, as did the car industry for the first seventy years of its existence. However the liability of vendors for dangerous design defects in cars and other physical products is now established beyond question. People who have been harmed by a defective product can sue the manufacturer. If the harm is done to the person who bought the goods, then the contract of sale may try to exclude them from claiming, so 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). Thus an end-user license agreement cannot invalidate a European resident's right to claim damages, and neither does a third party (such as a hacker) tampering with the device if the design allowed a foreseeable attack.

A claimant generally has to establish causation, demonstrate that the harm could have been foreseen, and show that manufacturer did not discharge its duty of care [7]. 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. Similarly, if vulnerabilities are regularly found despite previous instances being fixed, then it is foreseeable that more are to come. For this reason it is clearly negligent to sell network-attached devices that cannot be patched; and there are now standards for vulnerability lifecycle management that we will discuss in the next section and that a vendor ignores at its peril.

Even so, general liability law needs a refresh. The fact that products work as parts of larger systems is a significant and growing issue for the Directive, especially since it does not cover services. Firmware in a physical device is very likely covered while the server software on which an IoT device relies could well count as a service. Thus if harm were caused by a defect in (or an attack on) GPS navigation, the vendor would be liable for

an embedded device like a Garmin Navigator but not for an online service such as Google Maps, even though they have the same function. As we move to a world in which physical devices routinely interact with online services, this needs to be tackled, or vendors will just put safety-critical functionality in the cloud to escape liability.

The EU's own Blue Guide does indeed set an ambitious scope:

> "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." [13]

So one of our recommendations is that the EU extends 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; in addition to 'software as a service', we're seeing 'security as a service' and 'network as a service'. If the effect of this trend is to blunt vendors' liability for products that kill, then the law will have to catch up.

3.2 Transparency

Whatever the mix of litigation and regulation, both private claimants and government regulators know a lot less about a defective product than the vendor does. The vendor has the accident history, 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 NTHSA 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 [26].

In the world of cybersecurity, transparency comes from breach disclosure laws and coordinated vulnerability disclosure. Breach disclosure has caused firms to take cybersecurity much more seriously than before; writing to 50 million customers is not cheap, so insurers start to take an interest too. Responsible vulnerability disclosure has evolved in the IT industry, and represents consensus on how to deal with a key aspect of security. Vulnerabilities are reported with a fixed confidentiality period during which the vendor can develop a fix, test it, and ship it to users as part of a regular upgrade cycle. This ensures that the vendor has a proper incentive to fix the problem, while minimising their own costs. The reporting is often done through a neutral third party, such as a CERT; and 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 now an ISO standard for vulnerability disclosure, which should inform discussions of this topic [14]. This may be more significant than many realise, since the 'CE' mark that vendors affix to goods sold in the EU attests that it conforms to relevant standards. Thus, for example, toy vendors attest that they are not using lead paint. 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. Now that we have a vulnerability disclosure standard, firms which sell vulnerable network-attached devices might perhaps be held to account if their products cannot be patched – as was the case with some of the CCTV cameras recruited in October 2016 into the Mirai botnet. But it would be preferable for this to be made explicit, whether by a legal precedent or by regulation.

The EU's closest response to US security breach disclosure laws (pending the arrival of the General Data Protection Regulation (2016/679) in 2018) 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 be a CERT. Things are still a mess, though, as Europe has as many definitions of critical national infrastructure as countries; different industries collect different data; and there are no standards for calculating costs, whether private or public [4]. Worst of all, although ENISA collects data, it has neither the duty to pass it on to affected consumers and regulators, nor the technical staff to advise sector regulators directly. As legislation proceeds through European institutions it has no incentive to fight for safety to include security and no real access to the process, being located a long way from Brussels in Heraklion, Crete. This is a real problem, to which we will return in section 5.

3.3 Data protection

A third inspiration for Europe-wide safety regulation comes from privacy. EU data protection legislation has provided a framework for the protection of personal privacy since the 1980's. Member States have privacy regulators who require that processing of personal information be done according to fair processing principles, while Brussels has the European Data Protection Supervisor and the Article 29 Working Party of representatives of national regulators. This machinery is already 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 protection regime is unrelenting.

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 we fail to do so, then the integrity of our existing safety regulation mechanisms will be undermined. This paper attempts to give a warning of both the likely problems, and the great opportunities.

3.4 Attack and vulnerability testing

A fourth generic approach 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 'penetration testing' by an independent lab as part of the pre-market approval process. This would be no bad thing and has been advocated for 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 client. 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 causes them extra work.

Where a vendor seeks security testing of a product in order to help sell it, he will look for the testing lab that will give him the easiest ride. To tackle this conflict of interest, governments established the Common Criteria, under which testing labs are licensed by a national authority. The resulting certifications are used for some components of systems in government and banking, but the process is expensive to apply to whole systems; the test results can only speak to vulnerabilities in the tested device in a particular context and usually in a particular configuration. It is often harder 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 examples, and a whole field of research into 'security composability' [6, 10, 22]. In such cases, is natural for the two vendors to each claim that the bug is the other's fault.

Adversarial testing should ideally be of whole systems and companies, so the tester can explore and exploit context around devices, their usage, configuration, and the impact of a given vulnerability, or combinations of them.

Apart from the fact that most firms lack an incentive to have such tests done regularly, there are methodological complexities. In most cases, the operational security team of a target company knows the test is going on, which sometimes skews results; or they can artificially constrain the test scope, so it doesn't simulate a realistic opponent. Better practice is *red teaming*, adapted from the military, where penetration testers are not pre-announced to the operations staff. A red team can use a realistic range of techniques, such as phone-based social engineering, email phishing, or even physical intrusion – thus testing not only the software, but the organisation's overall capability. Red teams usually get in, but firms tend not to publicise the fact.

This can lead to vulnerabilities going undisclosed and unfixed for longer than need be. As penetration testers move from one company to another, they may find new vulnerabilities. They often recommend reporting a vulnerability to the vendor, only to have the asset owner decline. A few weeks later they find the same vulnerability in another company; indeed, dozens of asset owners may be vulnerable in the same way, simply because none of them was prepared to speak up. ENISA's report on testing of Industrial Control System devices noted and highlighted this problem as far back as 2013 [11].

The next problem, and perhaps the most serious in the long term, is that products are becoming much less static. As security and safety vulnerabilities are patched, regulators will have to deal with a moving target.

In the case of automobiles, the type approval regime will eventually have to incorporate not just pre-market safety and security testing, but a software update mechanism whereby security patches can be shipped. We expect to see the same with medical devices, electrotechnical equipment and much else. In principle, the coordinated vulnerability disclosure norms and programs seen in the software industry show how to organise this and make it work; in practice, the process will involve updating standards, test procedures and more. Managing this ecosystem will not be straightforward, and we will discuss the likely problems later on, in section 4.

3.5 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 [18].

The economics of cybercrime are not zero-sum; attackers profit very differently from defenders. A 2012 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 may steal money directly, or subvert advertising websites as a distribution engine for malvertising or malware.

Defenders' losses are sometimes just the attackers' gains. But much more frequently, they are very different. A user whose PC is infected by malware may be tempted into buying a new machine rather than cleaning up the old one; the small gains to the attacker are dwarfed by the profit made by the shop, the PC maker, and vendors such as Microsoft (whose vulnerable software may therefore actually be increasing its sales).

Vendors may operate under competition, as an oligopoly, or a monopoly. Especially in the latter cases, response may be slow. Microsoft left Windows more vulnerable than need be until customer anger (and the threat of Apple and Linux at the margins of its business) forced action. The absence of effective product liability for most desktop software may have also played a role.

Insurance of cyber risks is complex. They 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 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 so that it is not enough to simply analyse the frequency and severity of attacks and evaluate mitigating products. Insurers 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 sum lost rather than the 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.

Risks to society have been demonstrated in incidents such as the December 2015 Ukraine power outage, or the May 2017 WannaCry worm which led to a number of UK hospitals suspending accident and emergency admissions. In extreme cases there might be measurable effects on the GDP of the affected state or an impact on global supply chains. 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. These effects here are orders of magnitude greater than the actual sums stolen by cybercriminals.

Thus the goals for the regulators may include:

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

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

4 What needs to be done in practice

The transition from safety being a matter of pre-market inspection to a matter of monthly software upgrades will be a severe shock to the regulatory system. But there is much on which we can build, from the existing network of safety standards, regulators and testing labs, to security standards and industry practices.

4.1 Coordinated disclosure and long-tail maintenance

We mentioned the ISO vulnerability disclosure and coordination standards ISO/IEC 29147 and ISO/IEC 30111. The underlying substance is the vulnerability disclosure norms and practices developed by the IT industry. Twenty years ago, vendors tried to hush up details of breaches, using NDAs, legal threats and public relations techniques. 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 cars, air-conditioners, industrial control system components and medical devices are still mostly stuck in the twentieth century.

The reality of keeping complex online systems secure will move equipment vendors, systems integrators and operators towards regular patching, and regulators should help move their industries to standard processes for reporting, disclosure and patching. Regulators need to think about the incentives within their particular market. The usual carrot is reduced liability for compliant systems; a stick might be increased liability for firms that sell vulnerable network-attached devices that they cannot repair remotely. Some of these incentives will arise anyway – remote patching of vulnerable cars will be a lot cheaper than having them returned to a garage. But regulators should engage in such changes.

They need to anticipate issues that the market might not fix on its own, such as patches for second-hand vehicles.

If someone is writing code today for a car that will be sold in 2020, who will make security patches available in 2030, 2040 and indeed 2050? Most two-year old phones don't get patched because the OEM and the mobile network operator can't get their act together. So how on earth are we going to patch a 25-year-old Land Rover that spent 10 years in the British countryside and was then exported to Kenya?

Could we not just force the car and component vendors to open-source their software after ten years? The problem is the cost of testing. Vehicle software is usually tested in a 'lab car' which contains not just all the CPUs found in a vehicle but equipment to simulate driving and record behaviour. These test rigs cost from $50,000 and car makers understandably want to rebuild them to test newer models, rather than tying up millions of dollars in capital by keeping hundreds of old lab cars in service. When even Google ends its own patch support for Nexus phones after three years, what can we expect of a car maker whose test environment costs so much more to maintain? We suspect that regulators will simply have to compel the vendors to maintain a patching capability for decades. As the embedded energy cost of a car is about equal to its lifetime fuel burn, any substantial reduction in car lifetime will have ominous consequences for carbon targets.

But there may be opportunities for cost sharing, for example by tying software patches to hardware spares. The cost of car insurance claims is already rising because of the need to replace sensors such as radars and lidars when cars are damaged in low-speed collisions. A modern wing mirror containing a camera for automatic lane keeping can cost $1,000 to replace, compared with $100 for a common motorised wing mirror and $10 for a metal wing mirror from the 1970s. Car part makers have an incentive to help maintain the systems that use their components and thus drive their sales.

There will inevitably be issues around the interaction between subsystems. It will not be enough for the vendors of a car's automatic emergency braking, automatic lane-keeping, stability control and so on to maintain and patch their systems independently. The coordination costs will be non-trivial leading to an inevitable role of a systems integrator, probably the car vendor. In order to reduce bloat and increase intrinsic sustainability, we will need incentives in this ecosystem to drive collaboration rather than blame-shifting, to incentivise minimalism in the tool chain, and to get stable interfaces between subsystems that are verifiably safe and secure. If the market is left to itself, we can expect that subsystem vendors will attempt to make interfaces more complex in order to exclude competitors and lock in vendors of complementary products. We note that the European Commission has some experience of dealing with such issues in the competition law cases against Microsoft. We will discuss whole-system regulation further in section 4.7 below.

4.2 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 lead many people to avoid banking or shopping online, which imposes real costs. One of

the roles of the data protection authorities, for example, is to provide enough reassurance about privacy to make citizens comfortable with online shopping.

Similarly, safety regulators in the Internet of Things will earn some of their living by providing the comfort needed for the uptake of new technologies. For this to be sustainable, the trust which users place in them must be well-founded. Yet public confidence in vehicle type-approval regulation has been shaken by the Volkswagen scandal, which showed that car vendors were vigorously gaming emission standards.

This strengthens the case for open, standardised vulnerability management across all sectors where security is becoming a critical component of safety. Patching must not just be done, but be seen to be done. Safety regulators should require their industries to adopt a secure software development lifecycle, with a documented vulnerability management process. And once systems become dynamic, with regular software upgrades, safety regulation must be dynamic too. It must deal not just with bugs found in code, but with changing environments and emergent threats.

4.3 Collecting and publishing data

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 product liability and type approval, but because of detailed accident data, which are increasingly public. Insurers use accident statistics to set premiums; local authorities use them to prioritise road improvements; car vendors tear down crashed vehicles to see how to make future vehicles more crashworthy.

An essential market signal for security is the rate at which flaws are found, and the speed with which they are fixed. Current policy, following the NIS Directive, leans towards holding vulnerability information closely between ENISA and Member State security agencies, The IT industry, on the other hand, follows norms of responsible disclosure whereby vulnerability data are available to all, after a time-limited delay while systems are patched. In the USA, the NSA review group set up by President Obama after the Snowden revelations recommended that the NSA disclose the great majority of flaws for repair, rather than keeping them to exploit. President Obama accepted the group's other recommendations but equivocated on this one. The WannaCry worm, which used a leaked NSA exploit of an SMB vulnerability in Windows, now looks like forcing a rethink of this 'equities issue', as the intelligence community calls it. European institutions should also engage in this rethink and establish responsible disclosure as policy.

The same holds for statistics on crime. Electronic car theft tools are already used to exploit remote key-entry systems and immobilisers; we can expect much more of this. Governments have a duty to collect decent statistics, so that all can understand what's happening and those stakeholders with the ability to fix things can do so.

4.4 A caveat

If a banking app on your mobile phone is hacked, do you blame the network operator, the handset maker, the firm that wrote the operating system, the app developer or the

bank? Regulators cannot be everywhere, and transactions involving intermediate goods are generally left to the market. Thus a bank regulator will typically insist that banks make fraud victims good, and leave it up to them how they deal with their suppliers.

This may change. There is a specific problem with cars in that vendors often cannot get suppliers either to guarantee long-tail software maintenance, or to give the vendor access to the source code to do it themselves. As a result some car components are black boxes that cannot be maintained at all. It may become necessary for regulators to intervene further up the supply chain. This may also be efficient for products used by multiple car makers, or of new and rapidly evolving systems such as autopilots. And, as we noted in section 4.1 above, it may be both fair and efficient for component makers to shoulder some of the cost burden of long-tail software maintenance.

4.5 Who investigates incidents, and who gets the data?

Imagine that a terrorist takes control of a driverless vehicle, and runs it through a street filled with pedestrians. Who will provide the expertise to support the police, and would an engineer employed by the manufacturer be entirely frank about the novelty of the hacking exploit that was used?

What other regulations will be called for? The police will want a way to disable rogue vehicles remotely, and insurers will want 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. If a child figures out how to hack the system and break the rules, that may also need to be reported and investigated.

In short, the definition of 'security' for everyday things is extremely complex and context-dependent – and it will evolve over time in response to incidents and accidents.

There are over 25,000 road traffic fatalities in the EU every year with fifty times that number of minor injuries. This is half the number of 15 years ago, and if autonomous vehicles can cut the number by half again, that will be a huge societal benefit. But to obtain this benefit we need the data so that systems can learn, and the volume will be so large that we will need automatic means of collecting and analysing it.

Similarly, Facebook, WhatsApp and old-fashioned email are involved in large numbers of serious criminal cases, as well as very large numbers of less serious incidents such as cyberbullying at school. The volume of demand for access to data has led the big service companies to set up large help centres; caused law enforcement to train many officers to act as contact points; and led many countries to overhaul laws governing access.

A thorny problem here is jurisdiction. It's annoying if a small-town police officer needs to go through a cumbersome Mutual Legal Assistance Treaty (MLAT) process to get data from the USA to investigate a cyberbullying case. Many countries are passing data localization laws, requiring service firms to keep data within the jurisdiction (or easily available from it). Will we see this in the IoT as well? Will an autonomous vehicle vendor have to keep logs on servers in over 50 different countries in Africa alone?

More generally, who will investigate adverse events? Which state bodies (police, safety regulators, security regulators) need to be capable of which safety, security, and privacy

investigations? How do we prevent manage the inevitable turf wars, such as between intelligence agencies who want covert 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. Ultimately each regulator will have to curate an ecosystem in which incident investigation and adverse event reporting feed back into the remediation of hazards and the evolution of standards to give a system with both the mechanisms for learning, and the incentives.

4.6 Software patches as optimised product recall

Since the 1970s, regulators have forced manufacturers to recall and fix dangerous products. Since the NHTSA started ordering vehicle recalls, the threat of such a recall has forced vendors to pay attention to quality and safety. A recall of millions of vehicles causes 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 steadfast in requiring hazards to be fixed, manufacturers will have an incentive to set up the machinery to support a patch cycle.

The key thing will be to keep up the pressure. Once a safety fix no longer costs several billion dollars for a recall, but only several million dollars for investigation, patch development, testing and shipment, then regulators must start requiring fixes for hazards that cause millions in damage rather than just for the hazards that justify interventions costing billions. This brings us to the question of how we manage evolving systems.

4.7 Maintenance and evolution of whole systems

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, professionals have started to play a role, along with regulators. You have to have your car inspected every year, and you need a garage with the appropriate software. In the EU, the Right To Repair Directive ensures that manufacturers don't lock customers in to their own dealer networks but make the necessary software available to everyone.

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 be bid for centrally by specialist firms to clean up malware [9], 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 lethal incompatibility. It will not be enough to certify components; we will have to test, certify and monitor whole systems. We already certify a whole car, not just its brakes and steering; and we accept 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 road safety authorities need the authority to look at the whole picture.

5 Cyber-covigilance: our proposal

This paper has described how Europe has a multitude of sectoral regulators and standards agencies, giving details of how safety is being handled for road transport, medical devices and energy. These regulatory ecosystems will be disrupted 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 many years. The combination of software 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. So are the industries themselves; until a few years ago, no car company executive had any idea that their firm needed to hire a cryptographer. But industry at least has made a start. In the case of cars, the Escar conference on electronic security in cars has been going for ten years, with the typical car maker or large component vendor sending a few engineers; such a firm might have a few knowledgeable engineers plus several dozen who have been sent on security courses. Much the same can be found in the power equipment sector, where a large vendor with 100,000 employees might have three or four security specialists among its 10,000 engineers.

What can we reasonably expect of a regulator with 100 staff, most of whom are economists or lawyers? It is already an issue that the typical medical device regulator in Europe employs no engineers (while the FDA in America apparently employs only two). It would be splendid if safety regulators hired engineers, but if they find the resources to do that then the priority is likely to be a safety engineer to advise their lawyers and economists. And quite apart from budgetary constraints, government agencies in many countries (not just the EU) find it difficult to hire and retain good scientific and engineering staff for two reasons. First, civil service salary scales often don't allow them to compete with engineering firms. Second, in some administrative cultures, scientists and engineers are made to feel like second-class citizens; they discover that they have no promotion prospects beyond the specialist post for which they were hired, and leave. This is already a problem for ENISA, which suffers from two further issues: its focus on the national-security aspects of cybersecurity rather than the civilian aspects of privacy, safety and consumer protection; and the fact that it is situated in Heraklion, Crete, which is simply too remote from Brussels.

So what are the practical politics? A thorough institutional reform would lead to Europe amalgamating regulators into fewer organisations with bigger budgets, which would be staffed by engineers as well as economists and lawyers. However in the current political circumstances a thorough reform is unlikely. Consolidating existing institutions is painful,

because of the political cost of closing the predecessor institutions, which their host Member States would normally veto. Achieving such a change would require concerted action by a supermajority of heads of government and the losers would have to be compensated somehow. In the current circumstances, it is unlikely in the extreme that reforming safety regulation could get high enough on the agenda for radical action to be feasible. However, creating a new institution is relatively straightforward, especially if the budget is initially modest. Once the new agency becomes part of the European institutions, it can plug itself into internal networks, make itself useful, and lobby for more budget and personnel like everybody else.

Of course it would be great if ENISA could somehow move from Heraklion to Leuven, replace half its staff with engineers, and turn its focus from the security and intelligence agencies to consumer protection; but those agencies would no doubt advise Member State governments to vote against that. Having observed the evolution of European institutions, it looks a lot easier to start with a clean sheet of paper.

In our research report, we therefore proposed the creation of a European Safety and Security Engineering Agency to provide a shared resource for policymakers and regulators. It should ideally be located in Brussels – or failing that in a European city with good transport links to Brussels and a substantial existing tech sector (Berlin, Amsterdam, Paris, Munich, Helsinki, Tallinn and even Leuven come to mind; the list is not exhaustive).

Its mission will be to:

- support the European Commission's policy work where technical security or cryptography issues are relevant
- support sectoral regulators in the EU 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 from 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

To support these missions the new European Security Engineering Agency would need to collect 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 to bona fide researchers, so that research results are repeatable.

The existing safety test labs may have an important role to play in post-market surveillance: as safety and security become intertwined, we need safety people 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 best way to embed adversarial thinking into the community of safety engineers and certification bodies. The alternative is that they will learn the hard way, and we will all share the costs.

6 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 must start thinking about malicious adversaries as well as about random chance; about deception as well as about unforced errors.

There will be a substantial net benefit to society; we might hope to cut more then ten thousand road traffic fatalities per year across Europe, a similar number of fatal accidents involving medical equipment, and hundreds of thousands of injuries. But to secure these gains, and to mitigate the bad and the ugly aspects of technology, we will need more work from our regulators. What's more, this work will be more complex than they are used to, 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 safety 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); but most of the hard work will go into the 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. For these reasons, a multi-stakeholder approach involving co-vigilance by multiple actors is inevitable.

A second big change is that, as the centre of gravity shifts from general regulation to the sector-specific type, cybersecurity will not be regulated so much by rows (data privacy, national security, liability, transparency) but by columns (cars, planes, medical devices, toys ...). In many sectors it will be about safety first, with privacy and other aspects of consumer protection being secondary goals.

The strategic political challenge facing the European Union is whether it wants to be the world's safety regulator. If it rises to this challenge, then just as engineers in Silicon Valley now consider Europe to be the world's privacy regulator, they will defer to Europe on safety too. The 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. We therefore recommend the establishment of a European Safety and Security Engineering Agency to provide security expertise to Brussels policymakers and sectoral regulators as they navigate this technological shift.

The strategic educational challenge is that as safety and security become intertwined, cultures and working practices will change. Safety engineers will have to learn adversarial thinking while security engineers will have to think more about usability and maintainability. This will be the work of decades. At Cambridge we have already reorganised our teaching so that first-year undergraduates get an introductory course in 'Software and security engineering' where security and safety are taught as two aspects of the same mission: designing systems to mitigate harm, whether caused by adversaries or not.

The strategic research challenge will include how we make systems more sustainable. At present, we have enough difficulty creating and shipping patches for two-year-old mobile phones. How will we continue to patch the vehicles we're designing today when they are 20 or 30 years old? How can we create toolchains, libraries, APIs and test environments that can be maintained not just for years but for decades? And how can this work within a learning system where we know that new hazards and vulnerabilities – and new types of hazard and vulnerability – will emerge, and we plan to monitor incidents and learn from them? But how can we have stable tools and test environments, when (in regulated industries) the safety standards evolve over time while (in unregulated ones) the nature of what hazards and vulnerabilities are 'reasonably foreseeable' evolves too?

Finally, if the security economics community comes to see itself as being in the business of dependability economics, what is the shape of this new subject? What do we need to learn about the economics of safety? What new theoretical models do we need? What sort of data do we need to collect? And what will be the mechanisms whereby we influence policy and practice in the real world?

Acknowledgements

We are grateful to the European Commission, and in particular Gianmarco Baldini of the EC Research Centre, for commissioning the research that underpins this paper, and for permission to publish an abridged account of our findings. We are also grateful to Mike Ellims and Graeme Jenkinson for feedback on vehicle software, as well as to Robert Watson and other colleagues in the Cambridge security group for discussions of security sustainability.

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] Matt Bishop: *What is computer security?* IEEE Security & Privacy, 1, 2003

[7] Michael G Bridge: *Benjamin's sale of goods*. Sweet & Maxwell, 2012

[8] Pilita Clark and Sam Jones: *GCHQ intervenes to secure smart meters against hackers*. (accessed February 27, 2017) Financial Times, `http://www.ft.com/cms/s/0/ca2d7684-ed15-11e5-bb79-2303682345c8.html` 2016

[9] Richard Clayton: *Might governments clean-up malware?* Communication and Strategies, 2011

[10] Anupam Datta, Ante Derek, John C Mitchell, Arnab Roy: *Protocol composition logic (PCL)*. Electronic Notes in Theoretical Computer Science, 2007

[11] ENISA: *Good Practices for an EU ICS Testing Coordination Capability*. 2013

[12] European Commission: *Revisions of Medical Device Directives*. (accessed February 27, 2017) `http://ec.europa.eu/growth/sectors/medical-devices/regulatory-framework/revision_en`, 2016

[13] European Commission: *The 'Blue Guide' on the implementation of EU product rules 2016*. 2016

[14] International Organization for Standardization: *ISO/IEC 29147:2014 Information technology – Security techniques – Vulnerability disclosure*. (accessed February 27, 2017) ISO, Geneva, Switzerland, `http://www.iso.org/iso/catalogue_detail.htm?csnumber=45170`

[15] 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

[16] Paolo Masci, Rimvydas Rukenas, 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

[17] Jerry L Mashaw, David L Harfst: *The struggle for auto safety*. Harvard University Press Cambridge, MA, 1990

[18] Steven J Murdoch, Mike Bond, Ross Anderson: *How Certification Systems Fail: Lessons from the Ware Report*. Computer Laboratory, University of Cambridge 2012

[19] Monte Reel and Jordan Robertson: *It's Way Too Easy to Hack the Hospital*. Bloomberg Businessweek, Nov 2015

[20] Karen Sandler, Lysandra Ohrstrom, Laura Moy, Robert McVay: *Killed by code: Software transparency in implantable medical devices*. Software Freedom Law Center, 2010

[21] Harold Thimbleby: *Improving safety in medical devices and systems*. IEEE International Conference on Healthcare Informatics (ICHI), IEEE, 2013

[22] Ken Thompson: *Reflections on trusting trust*. Communications of the ACM, 27(8), 1984

[23] Iain Thomson: *Our pacemakers are totally secure, says short-sold St Jude*. (accessed February 27, 2017) The Register, `http://www.theregister.co.uk/2016/08/29/st_jude_hits_back_at_shortselling_security_firms_claims`, 2016

[24] Dan Tofan, Theodoros Nikolakopoulos, Eleni Darra: *The cost of incidents affecting CIIs*. ENISA, 2016

[25] *US vs. Guidant LLC 2011*. 2011

[26] Wikipedia, The Free Encyclopedia: *Ford Pinto*. (accessed February 27, 2017) `https://en.wikipedia.org/wiki/Ford_Pinto`