

S₁ E₁ C₃ U₁ R₁ I₁ T₁ Y₄
E₁ C₃ O₁ N₁ O₁ M₃ I₁ C₃ S₁
A₁ N₁ D₂ T₁ H₄ E₁
I₁ N₁ T₁ E₁ R₁ N₁ A₁ L₁
M₃ A₁ R₁ K₅ E₁ T₁

Ross Anderson
Rainer Böhme
Richard Clayton
Tyler Moore

Disclaimer

In August 2007, the European Network and Information Security Agency (ENISA) tendered a study related to the overall subject matter of “Barriers and Incentives for network and information security (NIS) in the Internal Market for e-Communication.” Views and opinions expressed in this report do not necessarily reflect those of ENISA.

Security Economics and The Internal Market

1 Executive Summary

Network and information security are of significant and growing economic importance. The direct cost to Europe of protective measures and electronic fraud is measured in billions of euros; and growing public concerns about information security hinder the development of both markets and public services, giving rise to even greater indirect costs. For example, while we were writing this report, the UK government confessed to the loss of child-benefit records affecting 25 million citizens. Further revelations about losses of electronic medical information and of data on children have called into question plans for the development of e-health and other systems.

Information security is now a mainstream political issue, and can no longer be considered the sole purview of technologists. Fortunately, information security economics has recently become a live research topic: as well as collecting data on what fails and how, security economists have discovered that systems often fail not for some technical reason, but because the incentives were wrong. An appropriate regulatory framework is just as important for protecting economic and other activity online as it is offline.

This report sets out to draw, from both economic principles and empirical data, a set of recommendations about what information security issues should be handled at the Member State level and what issues may require harmonisation – or at least coordination. In this executive summary, we draw together fifteen key policy proposals. We held a consultative meeting in December 2007 which established that almost all of these proposals have wide stakeholder support. We believe they will provide a sound basis for future action by ENISA and the European Commission.

Recommendations

1: There has long been a shortage of hard data about information security failures, as many of the available statistics are not only poor but are collected by parties such as security vendors or law enforcement agencies that have a vested interest in under- or over-reporting. Crime statistics are problematic enough in the traditional world, but things are harder still online because of the novelty and the lack of transparency. For example, citizens who are the victims of fraud often have difficulty finding out who is to blame because the incidents that compromised their personal data may have been covered up by the responsible data controllers. These problems are now being tackled with some success in many US states with security-breach reporting laws, and Europe needs one too.

We recommend that the EU introduce a comprehensive security-breach notification law.

2: Our survey of the available statistics has led us to conclude that there are two particularly problematic ‘black holes’ where data are fragmentary or simply unavailable. These are banks and ISPs. On the banking side, only the UK publishes detailed figures for elec-

tronic fraud, broken down by the types of attack. Similar figures are probably available to regulators in other Member States but are not published.

We recommend that the Commission (or the European Central Bank) regulate to ensure the publication of robust loss statistics for electronic crime.

3: On the ISP front, it is widely known in the industry that well-run ISPs are diligent about identifying and quarantining infected machines, while badly-run ISPs are not.

We recommend that ENISA collect and publish data about the quantity of spam and other bad traffic emitted by European ISPs.

4: People who leave infected machines attached to the network, so that they can send spam, host phishing websites and distribute illegal content, are polluting the digital environment, and the options available are broadly similar to those with which governments fight environmental pollution (a tax on pollution, a cap-and-trade system, or private action). Rather than a heavyweight central scheme, we think that civil liability might be tried first, and suggest

We recommend that the European Union introduce a statutory scale of damages against ISPs that do not respond promptly to requests for the removal of compromised machines, coupled with a right for users to have disconnected machines reconnected if they assume full liability.

5: A contentious political issue is liability for defective software. The software industry has historically disclaimed liability for defects, as did the motor industry for the first sixty years of its existence. There have been many calls for governments to make software vendors liable for the harm done by shoddy products and, as our civilisation comes to depend more and more on software, we will have to tackle the ‘culture of impunity’ among software developers.

We take the pragmatic view that software liability is too large an issue to be dealt with in a single Directive, because of the large and growing variety of goods and services in which software plays a critical role. Our suggested strategy is that the Commission take a patient and staged approach. There are already some laws that impose liability regardless of contract terms (for example, for personal injury), and it seems prudent for the time being to leave standalone embedded products to be dealt with by regulations on safety, product liability and consumer rights. Networked systems, however, can cause harm to others, and the Commission should start to tackle this. A good starting point would be to require vendors to certify that their products are secure by default.

We recommend that the EU develop and enforce standards for network-connected equipment to be secure by default.

This need not mean Common-Criteria certification of consumer electronics; it would be quite sufficient for vendors to self-certify. However, the vendor should be liable if the certification later turns out to have been erroneous. Thus if a brand of TV set is widely compromised and becomes used for hosting phishing and pornography sites, the ISPs who paid penalty charges for providing network connectivity to these TV sets should be

able to sue the TV vendor. In this way the Commission can start to move to a more incentive-compatible regime, by relentlessly reallocating slices of liability in response to specific market failures.

6: There has been controversy about vulnerability disclosure and patching. Recent research has shown that the approach favoured by the US Computer Emergency Response Team (US CERT) – namely responsible disclosure – gets better results than nondisclosure or open disclosure. However, some firms still take a long time to issue patches for vulnerabilities, and we believe that liability would help them along.

We recommend that the EU adopt a combination of early responsible vulnerability disclosure and vendor liability for unpatched software to speed the patch-development cycle.

7: Vendors also dissuade people from patching by bundling patches with upgrades and with disfeatures such as digital rights management.

We recommend security patches be offered for free, and that patches be kept separate from feature updates.

Likely future steps include making end-users liable for infections if they turn off automated patching or otherwise undermine the secure defaults provided by vendors. A useful analogy is that it's the car maker's responsibility to provide seat belts, and the motorist's responsibility to use them.

8: The next set of issues concern consumer rights. At present, the ability of consumers to get redress when they are the victims of fraud varies considerably across Member States. This issue was fudged during the preparation of the Payment Services Directive but now needs to be brought back on to the agenda.

The European Union should harmonise procedures for the resolution of disputes between customers and payment service providers over electronic transactions.

9: Some companies use marketing techniques that break various EU laws and/or exploit various loopholes in ways that should be banned or that provide cover for criminal activity. We need to abolish the business exemption for spam, criminalise firms who buy botnet services through third parties, and criminalise firms that install spyware on consumer computers without full user consent and without providing easy uninstallation.

We recommend that the European Commission prepare a proposal for a Directive establishing coherent regime of proportionate and effective sanctions against abusive online marketers.

10: The flip side of this is consumer protection, which will over time become much more complex than just a matter of payment dispute resolution. We already have an Unfair Contract Terms Directive, but stakeholders have raised other issues as well. Consumer protection in the broad sense is too wide for this report but will need attention.

ENISA should conduct research, coordinated with other affected stakehold-

ers and the European Commission, to study what changes are needed to consumer-protection law as commerce moves online.

11: The IT industry has tended towards dominant suppliers. As systems become increasingly interconnected, a common vulnerability could trigger cascading failures. Diversity, then, can be a security issue as well as a competitive one.

We recommend that ENISA should advise the competition authorities whenever diversity has security implications.

12: As for critical national infrastructure, one particular problem is the lack of appropriate incentives to provide resilience in competitive network markets.

We recommend that ENISA sponsor research to better understand the effects of Internet exchange point (IXP) failures. We also recommend they work with telecomms regulators to insist on best practice in IXP peering resilience.

13: As well as providing the right incentives for vendors and service providers, and protection for consumers, it is important to catch cyber-criminals, who at present act with near impunity thanks to the fragmentation of law-enforcement efforts. In order for the police to prosecute the criminals they catch, cyber-crimes must be offences in all Member States.

We recommend that the European Commission put immediate pressure on the 15 EU Member States that have yet to ratify the Council of Europe Convention on Cybercrime.

14: Furthermore, as nearly all cyber-crimes cross national borders, cooperation across jurisdictions must be improved. Joint operations and mutual legal assistance treaties have so far proved inadequate.

We recommend the establishment of an EU-wide body charged with facilitating international co-operation on cyber crime, using NATO as a model.

15: Finally, a number of regulations introduced for other purposes have caused problems for information security researchers and vendors – most notably the dual-use regulation 1334/2000, which controls cryptography with a keylength in excess of 56 bits, and the implementations of the cybercrime convention in some Member States that have criminalised the possession of ‘hacking tools’ (which can also catch security researchers). The security industry needs a ‘friend at court’.

We recommend that ENISA champion the interests of the information security sector within the European Commission to ensure that regulations introduced for other purposes do not inadvertently harm security researchers and firms.

Contents

1	Executive Summary	3
2	Introduction	9
2.1	The online criminal revolution	9
2.2	Regulatory context	12
2.3	Security economics	13
2.4	Scope	15
3	Existing economic barriers to security	18
4	Information asymmetries	22
4.1	Security breach disclosure laws	22
	Recommendation 1: <i>Breach notification</i>	26
4.2	Metrics	26
4.2.1	What are the statistics for?	27
4.2.2	What statistics are already being collected?	28
4.2.3	Case studies of security statistics	29
4.2.4	How should statistics be collected?	33
4.2.5	Metrics derived from market price information	37
4.3	Information sharing	40
4.3.1	Costs and benefits of sharing	40
4.3.2	Examples of information sharing	43
4.4	Information sharing recommendations	44
	Recommendation 2: <i>Electronic crime statistics</i>	45
	Recommendation 3: <i>Bad traffic statistics</i>	46
5	Externalities	47
5.1	Fixing externalities using carrots	47
5.2	Fixing externalities using sticks	49
5.2.1	Control points	49
5.2.2	Policy options for coping with externalities	50
	Recommendation 4: <i>Removal of compromised machines</i>	54
6	Liability assignment	55
6.1	Analogy with car safety	55
6.2	Competition policy	56
6.3	Product liability	57
6.4	Software and systems liability options	59
	Recommendation 5: <i>Secure equipment by default</i>	60
6.5	Patching	61
6.5.1	Challenge 1: Speeding up patch development	62
	Recommendation 6: <i>Responsible disclosure and fast patching</i>	64
6.5.2	Challenge 2: Increasing patch uptake	64
	Recommendation 7: <i>Free and separate security patches</i>	65
6.6	Consumer policy	65

6.6.1	Fair contract terms	65
	Recommendation 8: <i>Electronic payment dispute resolution</i>	66
6.6.2	Protection against abusive practices	67
	Recommendation 9: <i>Sanction abusive online marketers</i>	68
6.6.3	Consumer protection in general	68
	Recommendation 10: <i>Consumer-protection law</i>	70
7	Dealing with the lack of diversity	71
7.1	Promoting logical diversity	71
	Recommendation 11: <i>Advise competition authorities</i>	73
7.2	Promoting physical diversity in critical national infrastructure	73
7.2.1	Common mode failures and single points of failure	73
7.2.2	Internet exchange points	74
7.2.3	Hacking the critical national infrastructure	76
7.2.4	Policy options	76
	Recommendation 12: <i>Study IXP failures</i>	77
8	Fragmentation of legislation and law enforcement	78
8.1	Criminal law	78
	Recommendation 13: <i>Ratification of Cybercrime Convention</i>	79
8.2	Improving co-operation across jurisdictions	79
8.2.1	Defining the problem	79
8.2.2	Methods for co-operation	80
	Recommendation 14: <i>EU-wide co-operation on cyber crime</i>	81
9	Other issues	82
9.1	Cyber-insurance	82
9.2	Security research and legislation	87
	Recommendation 15: <i>Champion information security research</i>	88
10	Conclusions	89
	List of Figures	90
	List of Tables	90
	List of Acronyms	91
	References	93
A	Information society indicators on security	104
B	Internet exchange points	111
C	Methodology	113

2 Introduction

Until the 1970s, network and information security were largely the concern of national governments. Intelligence agencies used eavesdropping and traffic analysis techniques against rival countries, largely in the context of the Cold War, and attempted to limit insofar as was practicable the penetration of their own countries' networks by rival agencies. A legacy of this period is that in many Member States, the national technical authority for information security is an intelligence agency (such as GCHQ/CESG in Britain). There are still national defence concerns entwined with information security, such as the protection of critical national infrastructure. As the Internet becomes fundamental to the provision of ever-more goods and services, a nation or region that suffered a prolonged loss or degradation of network service could face serious consequences. These would not just be the 'obvious' problems, such as the dependence on networks of emergency services and of critical services such as healthcare. Logistics are nowadays so automated that within a week or two deliveries of food to supermarkets might start to become erratic. However, as we will discuss in this report, the provision of resilience and assurance to critical infrastructures is no longer a problem that can be solved at the national level alone. Our countries have grown so interdependent that some action is also needed at the community level.

From the 1970s until about 2004, however, the centre of gravity in information security shifted from governments to companies. As firms became ever more dependent on networked computer systems, the prospect of frauds and failures has increasingly driven investment in research and development. (The EU market for add-on information security products and services amounts to some EUR 4.6 billion.) Although there has been much publicity given to incidents of 'hacking' in which outsiders – often bored juveniles – penetrated company systems, the real centre of gravity in corporate information security has been preventing abuse by insiders. In a well-run company, information security mechanisms are only one component of a much larger system of internal control and risk management. This system extends from technical access controls and audit trails through staff training and other cultural aspects to insurance and money-laundering controls; it aims largely at preventing frauds by the company's own staff. Of course there is some interaction between national-security and corporate concerns; economic espionage may target key companies as well as governments, and much of the critical infrastructure is now in private rather than public management. However the perspectives and incentives of private firms and public agencies are different.

2.1 The online criminal revolution

Since about 2004, volume crime has arrived on the scene. All of a sudden, criminals who were carrying out card fraud and attacks on electronic banking got organised, thanks to a small number of criminal organisations and a number of chat-rooms and other electronic fora where criminals can trade stolen card and bank account data, hacking tools and other services. Previously, a card fraudster had to run a vertically-integrated business: he might, for example, buy a card encoding machine, then get a job in a shop where he could take extra copies of customers' bank cards, and go out at night to draw cash from ATMs using card clones. Similarly, an electronic banking fraud might involve a corrupt

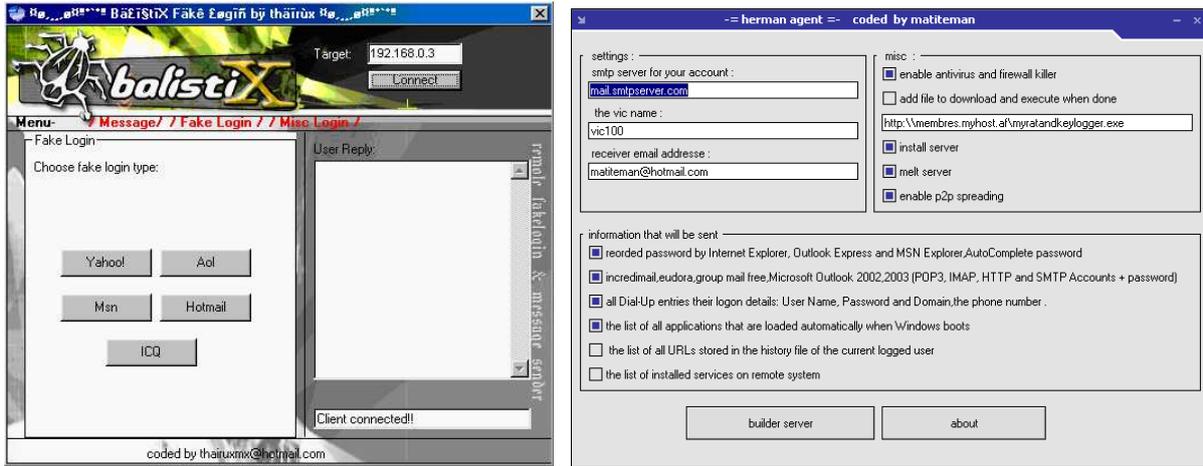


Figure 1: Web trojan generator interface (left) and data theft crimeware interface (right). Source: [43]

bank employee at a call center collecting password data for use by an accomplice. Such crimes were local and inefficient.

The emergence of criminal networks has changed that. Someone who can collect electronic banking passwords, or bank card and PIN data, can sell them online to anonymous brokers; and brokers sell them on to *cashiers* who specialise in money laundering. The money-laundering step becomes further specialised, with spammers recruiting *mules* who are duped into accepting bank payments and sending them onwards to third countries via Western Union. The collection of bank passwords has become further specialised as *phishermen* operate websites that appear to be genuine bank websites, and hire the spammers to drive bank customers to them. Both the spammers and the phishermen use malware writers, who create the hacking tools that compromise millions of machines. A new profession, the *botnet herder*, has arisen – the man who manages a large collection of compromised PCs and rents them out to the spammers and phishermen. On occasion, botnets can be used for even more sinister purposes, such as by blackmailers who threaten to take down bookmakers’ websites just before large sporting events – and, in the case of Estonia, to attack a Member State’s infrastructure as a political protest.

In the eighteenth century, rapid specialisation by artisans led to the Industrial Revolution. Adam Smith describes how a pin factory became more efficient by having one worker cutting the wire, another sharpening the pins, and so on; the last few years have seen an online criminal revolution driven along very similar lines.

Hacking has turned from a sport into a business, and its tools are becoming increasingly commoditised. There has been an explosion of crimeware – malicious software used to perpetrate a variety of online crimes. Crimeware used to require skill to create, but now it’s available almost as a consumer product. Keyloggers, data theft tools and even phishing sites can be constructed using toolkits complete with sophisticated graphical user interfaces. Figure 1 gives screenshots from two such tools. On the left is a web Trojan generator, which creates fake login pages for Yahoo!, AOL, Hotmail and others to be automatically overlaid on the authentic login pages. On the right is a tool for automatically scraping sensitive data from infected computers, such as the Internet Ex-

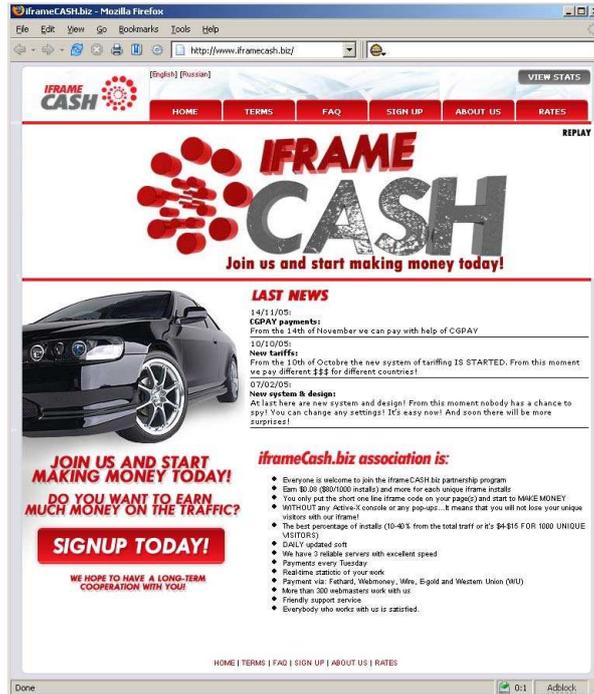


Figure 2: Crimeware affiliate marketing. Source: [43]

plorer saved password file and browsing history, along with the user’s email login details and loaded programs. The ‘quality’ of these tools is improving rapidly, as their authors invest in proper research, development, quality control and customer service. Most tools are not initially detected by the common antivirus products, as their authors test them against these products; and when the antivirus vendors do catch up, the crimeware authors issue updates. This is driving an escalating arms race of online attack and defence. (And volume crime facilitates both corporate and national-security crimes as it creates a background of general attack traffic within which criminals can hide, and also makes high-quality crimeware tools both widely available and easily usable.)

Most commonly, crimeware is spread by tricking users into downloading attachments from an email or a malicious web site. The attachments purport to be salacious photos, games, or even spam blockers. Symantec estimates that 46 % of malicious code propagated via email in the first half of 2007 [130]. Another option for spreading malware is to use exploits – Symantec also found that 18 % of the malware they examined exploited vulnerabilities. Most worrying, however, is that the distribution of crimeware is becoming more sophisticated as the criminal economy develops. For example, so-called affiliate marketing programs have been set up that pay web site operators to install crimeware on its visitors’ computers using exploits. Figure 2 shows a screenshot for one such affiliate marketing web site, which asks webmasters to install iframes pointing to an attacker’s site for installing crimeware. In return, the webmaster receives a commission ranging from USD 0.08 to USD 0.50 per infection [43].

2.2 Regulatory context

In May 2007 the European Commission issued a Communication ‘towards a general policy on the fight against cyber crime’ [47]. It noted that there is not even an agreed definition of cyber-crime, and it proposed a threefold definition:

1. traditional forms of crime such as fraud or forgery, though committed over electronic communication networks and information systems;
2. the publication of illegal content over electronic media (i.a. child sexual abuse material or incitement to racial hatred);
3. crimes unique to electronic networks, i.e. attacks against information systems, denial of service and hacking.

It also identified eight problem areas:

1. A growing vulnerability to cyber crime risks for society, business and citizens;
2. An increased frequency and sophistication of cyber crime offences;
3. A lack of a coherent EU-level policy and legislation for the fight against cyber crime;
4. Specific difficulties in operational law enforcement co-operation regarding cyber crime, due to the cross-border character of this type of crime, the potential great distance between the crime perpetrator and the crime victim and the extreme speed with which crimes can be committed;
5. need to develop competence and technical tools (training and research);
6. The lack of a functional structure for co-operation between important stakeholders in the public and the private sector;
7. Unclear system of responsibilities and liabilities for the security of applications as well as for computer soft- and hardware;
8. The lack of awareness among consumers and others of the risks emanating from cyber crime.

A number of EU Directives have set out the general framework for regulating the Internet. There are a set of five directives dating from 2002 (Access [54], Authorisation [55], Framework [56], Universal Service [57], and Privacy [58]) which regulate the telecommunications companies. They are currently under review, and proposals for their revision were published in November 2007 [46].

Consumer protection is provided by the 1997 Distance Selling Directive [51] and the 2000 E-Commerce Directive [53]. Additionally e-commerce is assisted (at least in principle¹) by the 1999 Electronic Signatures Directive [52]. Assistance to law enforcement is provided by the 2006 Data Retention Directive [60].

¹Because regulations are different in different jurisdictions (private keys must be escrowed, private keys must never be escrowed, etc) it has been found to be simpler to develop public key infrastructures by using contract law rather than digital signature law [1]

2.3 Security economics

The contribution of this report lies in the field of the economics of information security. We are focused largely on the third of the Commission's three types of cyber-crime, namely on the new offences involving attacks on information systems, denial of service and hacking, although network insecurity spills over into the other two categories as well. Our work has implications for most of the problem areas: our key message is that in order to solve the first two problems (growing vulnerability and increasing crime) the Commission must pay attention to the third and seventh (policy and legislation must coherently allocate responsibilities and liabilities so that the parties in a position to fix problems have an incentive to do so).

Network and information security is of growing economic importance in Europe (and elsewhere): sales of anti-virus software, cryptographic products, and services ranging from spam filtering through phishing-site 'take-down' to brand protection and copyright enforcement are in the billions of euros per annum. Add-on security products alone, such as anti-virus software, were estimated by Forrester to be worth an estimated EUR 4.6 billion in 2008, while our industry sources suggest that the market for financial-sector security products is EUR 1.5 billion. In addition, insecurity – and the perception of insecurity – has a significant impact in wider markets. Some people buy premium products, such as Apple computers, in the expectation that they will be less vulnerable to malware; and, as can be seen from Table 12 in the Appendix, a significant and growing number of people have failed to order goods or services over the Internet because of security or privacy concerns (in three countries – Germany, Finland and Cyprus – a majority of respondents were in this camp.) It thus appears that the indirect costs of Internet insecurity are billions of euros more.

Security economics research

The economics of security play a deeper role too. Since about 2000, researchers have realised that many security failures have economic causes. Systems often fail because the organisations that defend them do not bear the full costs of failure. For example, in countries with lax banking regulation, banks can pass more of the cost of fraud on to customers and merchants, which undermines their own incentive to protect payment systems properly. This led to a UK parliamentary committee recommending tighter bank regulation as one of the needed remedies for Internet insecurity.

In addition, so long as anti-virus software is left to individuals to purchase and install, there may be a less than optimal level of protection – as infected machines typically cause trouble for other machines rather than their owners. This has led to lobbying from the anti-virus industry for the purchase of their products to become compulsory. How is the legislator to assess such claims?

In addition, information security mechanisms are increasingly used to support business models. The best-known examples are the use of digital rights management (DRM) systems to regulate the use of music and film downloads, and the use of cryptographic authentication mechanisms in product tying – as when printers are designed to only work with ink cartridges from the same manufacturer, or video-games consoles are subsidised from sales of games software. Although such mechanisms can be economically efficient,

they are often unpopular, have side-effects, and may raise competition policy issues.

The shortage of data

The economic study of information security products and services is thus of rapidly growing relevance to policy makers, yet it has been troubled from its earliest days by the lack of a solid evidence base. For at least two decades, both governments and security vendors have been complaining about inadequate information security expenditure by companies, and have repeatedly suggested that firms such as banks under-report computer security incidents in order to avoid loss of confidence. Other observers have suggested that companies over-report the value of incidents in order to get the police interested in investigating them. The insurance markets are of some assistance in risk assessment, but not much – markets for cyber-risk cover were disrupted around the year 2000 by fears about the Millennium Bug, were not particularly competitive before then, and have not been completely satisfactory since. The recent introduction of security breach disclosure laws in many US states has gone some way towards filling the information gap, and studies into the effects of breach disclosures on company stock prices have also helped.

Over- and under-reporting can lead directly to incorrect policy choices. For instance, the number of phishing websites and distinct attackers has been consistently over-reported, suggesting that the problem is too large and diffuse for the police, despite the fact that only a relatively small number of players are behind the majority of attacks. While bank fraud in the English-speaking world is dominated by fake websites, in Continental Europe the main problem comes from keyloggers and session hijacking. The public is told that they should buy anti-virus software, but this is becoming ineffective as the malware writers become more professional and test their offensive products properly against the existing defensive products before releasing them. In fact the socially optimal response may now be a police response. The same may go for spam; while a few years ago spam may have been sent by large numbers of small firms, there is now evidence of consolidation, with most spam by volume being sent by the operators of a small number of large botnets.

Cross-border dimension

An important question is whether enforcement is likely to require action on a European rather than national scale. Since many attacks are global in scope, the impact of the attack in any one jurisdiction may not justify intervention, even when the overall impact justifies it. For example, the London Metropolitan Police might take the view that only 5% of phishing victims are from the UK, and maybe 1% are from London, so why should they expend effort in trying to catch a large Russian phishing gang? Yet a European agency may take the view that 30–40% of the victims are European, so European action is justified. The nature of the action is also an important question. In some cases, the EU can facilitate coordination between national police forces; in the case of the large Russian gangs, the EU might help the US authorities to bring diplomatic pressure on Moscow to close the gangs down. It might also help by providing rewards for information leading to the arrest and conviction of the individuals controlling particular criminal operations, or

in coordinating the provision of such rewards by banks and other victims.

Policy options

A number of information security problems can be solved by private action, but not all. Many institutions may struggle to see why they should co-operate by sharing attack data that could not just reveal technical weaknesses but expose them to litigation. This has led, in the USA, to public-sector information sharing initiatives, and also to private-sector companies that buy, broker or aggregate vulnerability information. In addition, vulnerabilities in one firm may result in claims against another: a compromised ATM operated by one bank may result in other banks receiving claims from customers whose cards have been cloned. Where banks can deny liability – as in the UK and Germany – this can undermine the incentive to co-operate. A quite different pattern is found with online fraud and phishing: in the UK, for example, one bank suffered some GBP 34 million of the GBP 36 million of total phishing losses in 2006, which eroded the incentives of all the other banks to co-operate. Thus, for a variety of reasons, the state will have a role to play, either as policeman, or regulator, or coordinator. The state can also act more subtly, for example by security-breach disclosure laws.

In the specific case of the European Union, regulatory options range from direct legislation (previous examples being the Data Protection Directive and the Electronic Commerce Directive), sector-specific regulation (such as the recent Payment Services Directive), coordinating groups (such as the Article 29 Working Party on data protection law), the funding of research, down to the collection and publication of information. Unfortunately, regulatory actions are subject to multiple political and lobbying forces that pull in different directions. As the May 2007 Communication makes clear, the EU needs to make its policy on information security more coherent and to ensure that it's taken into account when policy on related matters is being formulated.

2.4 Scope

Network and information security has huge and growing scope. As more and more devices acquire processors and communications, we move to world of 'pervasive computing' in which we will each have hundreds of computers embedded invisibly in our homes, cars and places of work. Already a high-end motor car has over 40 microprocessors in it. Security is an issue (can the engine control unit be modified by the driver to give higher performance? The vendor wants to stop this to prevent increased warranty claims) and spills over into policy (the vendor is required to make the unit tamper-resistant to prevent increased exhaust emissions under Directive 98/68/EC section 5.1.4).

There are dozens of other embedded systems where security and policy already meet, and as time goes on, most areas of government regulation are likely to experience information security issues. This presents us with a problem of scope and focus for this report. Following discussions with ENISA, we focus on the direct and systemic security threats to networked information systems consisting largely of network-connected computers, whether clients or servers; to the routers and other underlying communications infrastructure; and to services delivered to mobile phones, PDAs and other peripatetic devices. Embedded systems, whether in vehicles, in buildings, or worn on the person, do of course interact with core systems, and we will mention them in passing. However, full

consideration of their policy implications must be left to further reports.

We are also largely excluding from this report any discussion of government information-security systems; although they are historically important and are converging with the protection mechanism used in business and commerce, their classified background and their entanglement with defence procurement makes them too complex and distracting to be considered systematically here.

In this report we differentiate software developers into those who sell software to satisfy a demand for functional properties, e.g. operating systems, middleware and applications, and developers who complement these products with products adding other (often non-functional) properties such as security. We will consider the *security industry* to be in the latter category and to comprise vendors of anti-virus, firewall, intrusion-detection, anti-spam and anti-phishing technology. Of course there is some overlap: Microsoft owns an anti-virus software vendor while it also supplies anti-spyware products free of charge. However, for practical purposes, we need to draw a distinction. This is because in an ideal world, operating systems, protocols, and applications would be secure in the first place, so the multi-billion-euro security industry would be obsolete.

When it comes to the financial sector, we will consider the security products to be the cryptographic devices, fraud-detection software and other core security products. As remarked above, this will add some EUR 1.5 billion to a core security market of EUR 4.6 billion. There is some overlap between these figures as banks also buy anti-virus software as well as specialist systems for fraud detection and so on. (We exclude consumables such as bank cards: if they were included, a further EUR 500 million per annum would be added to the total). The reason for including the financial sector explicitly in this report is that cyber-crimes mostly affect citizens via financial fraud. Citizens do of course have other concerns, such as privacy; but the main perceived problem at the end of 2007 in most European countries is fraud. This takes a number of forms, from online credit-card fraud though bank account takeover as a result of keyloggers or phishing.

We thus consider the following fraud lifecycle.

1. **Design flaw:** A vulnerability may be introduced into a system during the design process, as with the vulnerabilities in the EMV payment card protocols.
2. **Implementation flaw:** A vulnerability may alternatively be introduced by careless implementation, as when programmers fail to check the length of input strings leading to buffer-overflow exposures.
3. **Vulnerability discovery:** An exploitable flaw is discovered. The discoverer may be a responsible researcher who reports it to the vendor, or an attacker who uses it directly (a *zero-day exploit*).
4. **Patching:** The vendor patches the exploit. In the case of an online service such as Google, a software change on the server can be done at once; in the case of an operating system it typically means shipping a monthly product update.
5. **Post-patch exploit:** The majority of exploits involve flaws for which patches are available, but on machines whose owners haven't patched them. Many users don't patch quickly (or at all) and many attackers reverse-engineer patches to discover the flaws that they were designed to fix.

6. **Botnet recruitment:** Many exploited machines are recruited to *botnets*, networks of machines under the control of criminals that are used for criminal purposes (sending spam, hosting phishing websites, doing denial-of-service attacks, etc).
7. **Bot discovery and disinfection:** Infected machines are identified (because they are sending spam, hosting illegal websites etc.) and the ISP (if following best practice) then takes them offline.
8. **Asset tracing and recovery:** Where criminals have succeeded in taking over a citizen's bank account and start to transfer money out, typically to 'mules' who launder it, the banks' fraud-detection systems notice this and freeze the account.

A proper policy analysis of cyber-crime needs to consider all these steps. System vendors make socially suboptimal protection decisions because of wrong incentives: security isn't free, and they will provide less of it than they should if privacy laws aren't enforced properly, or the costs of fraud fall on others. Ensuring that an adequate amount of security research gets done, and that most vulnerabilities are reported responsibly to vendors rather than sold to criminals, is also a matter of (sometimes complex) incentives. Patching introduces further tensions: an operating-system vendor might like to patch frequently, but as patches can break application software, this would impose excessive costs on other stakeholders (including customers who write their own application software). It would be ideal if users who don't maintain their own software patched quickly, but often security fixes are bundled with upgrades that many customers don't want.

Botnet recruitment would be much harder if popular applications such as browsers had more usable security; yet many of the existing mechanisms appear designed by techies for techies, which raises a number of liability and even discrimination issues. Many machines get infected when users click on links in email, and thus ideally payment service providers would not train their customers to click on links in emails; yet many do. And once machines are infected, it's good practice for ISPs to spot them and take them offline, by transferring them to a 'walled garden' from which their users can access anti-virus software but not do much else. But many ISPs don't do this, and as a result some ISPs are the source of much more malicious traffic than others. Finally, banks vary enormously in their capability at detecting fraud and dealing with it.

So market failures are involved in every step of the cyber crime process, and many of them have implications for the Single Market. We will now consider them by failure type – information asymmetries, externalities, incorrect liability allocation, monopoly/oligopoly, and fragmentation of legislation and law enforcement.

3 Existing economic barriers to security

We use the following framework to classify and analyse the economic barriers to network and information security in the subsequent sections.

1. Information asymmetries
2. Externalities
3. Liability dumping prompted by network convergence and interdependence
4. Lack of diversity in platforms and networks
5. Fragmentation of legislation and law enforcement

Information Asymmetries Asymmetric information – where one party to a transaction has better information than the other – can be a strong impediment to effective security. The study of this subject was initiated by George Akerlof’s Nobel-Prize-winning paper on the ‘market for lemons’, in which he imagined a town with 50 good used cars for sale (worth \$2000 each), along with 50 ‘lemons’ (worth \$1000 each). The sellers know the difference but the buyers do not, with the result that the market price ends up near \$1000 [3]. A lemons market also affects some security products and services, as their effectiveness is difficult for consumers to ascertain correctly. The consumers refuse to pay a premium for quality they cannot assess, so products and services tend to be of poor quality.

The tendency of bad security products to drive out good ones from the marketplace has long been known, and at present the main initiative supported by the Commission and Member State governments is the Common Criteria – a framework for product evaluation that evolved mostly for government-sector suppliers but is now being used as well by (for example) vendors of point-of-sale terminals. This is at least a start, but it has had little impact so far outside the government and (to a lesser extent) financial sectors. The public has inadequate information about the relative effectiveness of the many security products and services on general offer. It has also long been known that we simply do not have good statistics on online crime.

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. We assess these options later.

Externalities The effects (positive or negative) that economic transactions have on third parties are called externalities. Familiar examples are the industrial spin-off from scientific research (a positive externality) and environmental pollution (a negative externality).

Many important security threats are characterised by negative externalities. For example, home computers are increasingly being compromised and loaded with malware used to harm others (by sending spam, hosting phishing sites or launching denial-of-service attacks). The malware typically does not harm the user directly; it may even patch the user's computer, to prevent it being infected with competing malware! As a result, a user who connects an unpatched computer to the Internet does not face the full economic consequences of her action. For this reason, internet insecurity has been likened to air pollution: connecting an infected PC to the Internet is analogous to burning a smoky coal fire.

However, the analogy has its limits, and a case can be made that the average consumer isn't competent to detect and deal with infection. The consumer's ISP is in a much better position to detect infected machines, and to insist that they be cleaned up as a condition of continued service. Here a further set of externalities come into play. Small-to-medium ISPs have an incentive to clean up user machines (as being a source of spam would otherwise damage their ability to have their email accepted [123]) while large ISPs at present enjoy a certain impunity. We will consider several policy remedies for reducing the digital pollution emanating from ISPs, from taxation to a cap-and-trade system to fixed penalty charges.

Security investment can thus create quite complex externalities. Another example is that the benefit of a protective measure often depends on the number of users adopting it (a *network externality*). For example, encryption software needs to be present at both ends of a communication in order to protect it, and so the first company to buy encryption software can protect communications with its branches, but not with its customers or its suppliers. As a result, the cost of a new product or service may be greater than the benefit until a certain threshold number of firms adopt. Thus security products and services can be difficult to launch unless early-adopter firms can obtain sufficient benefits directly. Yet another example is that investments can be strategic complements: an individual taking protective measures may also protect others, inviting them to free-ride. Policy tools for overcoming such externalities range from standardisation through regulation and subsidy to strategic procurement.

Liability dumping A further bundle of problems relate to liability dumping. Firms seeking to manage risk often do so by externalising it on less powerful suppliers or customers. The most obvious example is the way in which software and service suppliers impose 'shrink-wrap' or 'click-wrap' licenses on customers disclaiming all liability, including for security failures, and in some cases also taking 'consent' to the installation of spyware. This is a public policy issue as it removes a major incentive for the emergence of a market for more secure languages and tools, and for the employment of professional software engineering methods. Yet a single vulnerability can lead to millions of euros of damage.

Another example is the problem of mobile phone security; mobile phones have a long and complex supply chain, starting from the intellectual property owners, the chipmaker, the software supplier, the handset vendor, the network operator and brand from which

the customer buys service. Each of these players seeks to have others bear the costs of security as much as possible, while using security mechanisms to maximise its own power in the chain. One side-effect has been the failure of the OMA DRM Architecture V2 to come into widespread use, which in turn is said to have depressed the market for music downloads to mobile phones.

A third example is in payment services. The recent Payment Services Directive goes some way towards harmonisation of service rules across the EU but still leaves consumer protection significantly behind the USA. Banks are allowed to set dispute resolution procedures by their terms and conditions, and do so in their favour – as found for example in the recent report of the UK House of Lords Science and Technology Committee into Personal Internet Security [76], which recommended that the traditional consumer protection enshrined in banking law since the nineteenth century should be extended to electronic transactions too. At the professional level, there is a concern that European SMEs cannot always get certain banking services necessary for e-business (and in particular the acquisition of credit card transactions) on terms comparable to their US competitors. This places European e-business at a disadvantage.

Lack of diversity Lack of diversity is a common complaint against platform vendors, whether Microsoft or Cisco or even Symbian. This is not just a matter for the European Commission Directorate General for Competition (DG COMP); lack of diversity makes successful attacks more devastating and harder to insure against. Homogeneous architectures share common vulnerabilities, and this increases the variance of the loss distribution due to security incidents. Such high variance undermines many firms' confidence in technology and makes them reluctant to invest.

One possible device for risk sharing and control is insurance; but high loss correlation renders large market segments uninsurable. Thus the market structure of the IT industry is a significant factor in society's ability to manage and absorb cyber risks, and has a negative effect on the markets for cyber-insurance.

Communication service providers are also affected; smaller ISPs find it cheaper to use single peering points, with the result that only large ISPs offer their customers resilience against peering point outage. This not only places these smaller ISPs (which are mainly SMEs and providing services to SMEs) at a disadvantage but shades over into critical national infrastructure concerns.

Fragmentation of legislation and law enforcement The fragmentation of legislation and law-enforcement jurisdictions hinders rapid response. Mitigating many attacks requires better and faster co-operation across jurisdictions. For example, the most important factor in deterring and frustrating phishing attacks is the speed of asset recovery. A bank learning of a customer account compromise needs to be able to trace and freeze any stolen assets quickly. The phishermen for their part use offshore money transfer services and, as these are shut down, they are increasingly sending hot money through the banks of Member States with a relaxed attitude to asset recovery. This issue is also of interest to authorities tackling money laundering, and spills over from first pillar to third pillar issues, but the proper functioning of the Internal Market also depends on enforcement tasks that stop short of police involvement. An example is the enforcement of trading standards, which in the UK is largely the domain of county councils; these bodies are largely set up

to inspect local traders, and lack the expertise to tackle complaints of online scams. The question thus arises of whether we need a European Trading Standards Agency. Another example is that the Single Market also requires predictably dependable payments and public trust in payment service providers – which cannot nowadays be divorced from NIS.

Some security problems are a mixture of the above. To take a concrete example, in October 2007 a ‘skimmer’ was found on a cash machine in St Andrews Street, Cambridge, and the police duly alerted the public that if they had used that machine they should check their bank statements and call their bank if there was any fraud [11]. (Skimmers are devices attached to the ATM card slot that copy the magnetic-strip data and contain a small camera to record the PIN; they are available online for USD 500. They are used to make magnetic-strip copies of debit cards that are then used in ATMs that allow magnetic-strip fallback). Thus local bank customers who heard the news on local radio were in a position to complain and have their losses made good. However if a businessman visiting Cambridge from Germany had used that cash machine, and gone home the following day, then in all probability he would not have heard the news, and when he complained to his bank in Germany about unauthorised transactions he would most likely have been told that since his card had a chip in it, and the PIN was used, he was liable.

Such failures are a mixture of asymmetric information (the bank knows more about the risks than either merchants or customers), liability dumping (banks in the UK and Germany have been particularly successful at dumping the risks of fraud on their customers) and fragmentation of legislation and law enforcement. They are clearly a single-market issue, as current procedures discriminate against non-local customers. Fortunately, there are fairly straightforward ways to deal with such failures – such as security breach disclosure laws, which we shall discuss in the next section.

In the following sections, we take each of these barriers in turn, discussing available solutions and recommending the best course of action.

4 Information asymmetries

In this section we describe ways to reduce information asymmetries. There is a growing consensus, among not just stakeholders but the wider policy community, that fixing information asymmetries requires a breach disclosure law as outlined in Section 4.1. It not only makes gathering statistics easier, but also empowers victims to get redress and take precautions, while shaming lazy companies into taking action. In Section 4.2, we discuss other available data sources and requirements for robust security statistical indicators. In Section 4.3, we outline conditions for stakeholders to share relevant data and make recommendations to increase data-sharing.

4.1 Security breach disclosure laws

The first ‘security breach disclosure’ law to be enacted in the United States was California’s A.B.700 in September 2002 [19], which came into force as Cal. Civil Code §1798.29, in July 2003. It applies to public and private entities that conduct business in California and requires them to notify affected individuals if personal data under their control is believed to have been acquired by an unauthorised person. The definition of personal data is restricted to a name combined with a Social Security Number, a driver’s license number, or credit/debit card number along with a password. In practice most computerised records holding what a European would call ‘personal data’ are likely to be covered. If personal data is ‘lost’ then the entity is obliged to inform the people who are affected.

The intention of the law was twofold. It was intended to ensure that when data was found to have been stolen, individuals would have the opportunity to take appropriate steps to protect their interests – such as putting a ‘lock’ on their file at credit agencies. It was also intended to provide an incentive on companies holding personal data to take steps to keep it secure. In particular, the law makes it clear that if data is encrypted, then in most circumstances it would not be deemed to have been lost, even if someone unauthorised obtained the encrypted material. This might be expected to promote the use of encryption to protect personal data.

Initially there was considerable publicity when companies lost data and people were informed, but the novelty quickly faded, and only very large or unusual security breaches make it into the media. In some cases, such as the ChoicePoint scandal where criminals were able to access 163,000 credit reports, there has been a substantial impact on the stock price – not least because the regulator subsequently fined the company USD 15 million [65]. Acquisti et al. [2] have studied this issue and found that there is a statistically significant negative impact on stock prices.

The ChoicePoint case and some other high-profile security breaches led to the California law being followed by further laws in at least 34 other states [106], although they differ somewhat in their details. In particular some of them permit companies to assess the risk and they need not issue a notification if they believe there is ‘no risk’. Some of the state laws require that their citizens be notified ‘first’ which is difficult for companies with a national presence! The variations between the laws has led to calls for a federal statute, but although bills have been introduced in Congress, none have had much success so far.

The Privacy Rights Clearinghouse publishes a database of known security breaches and

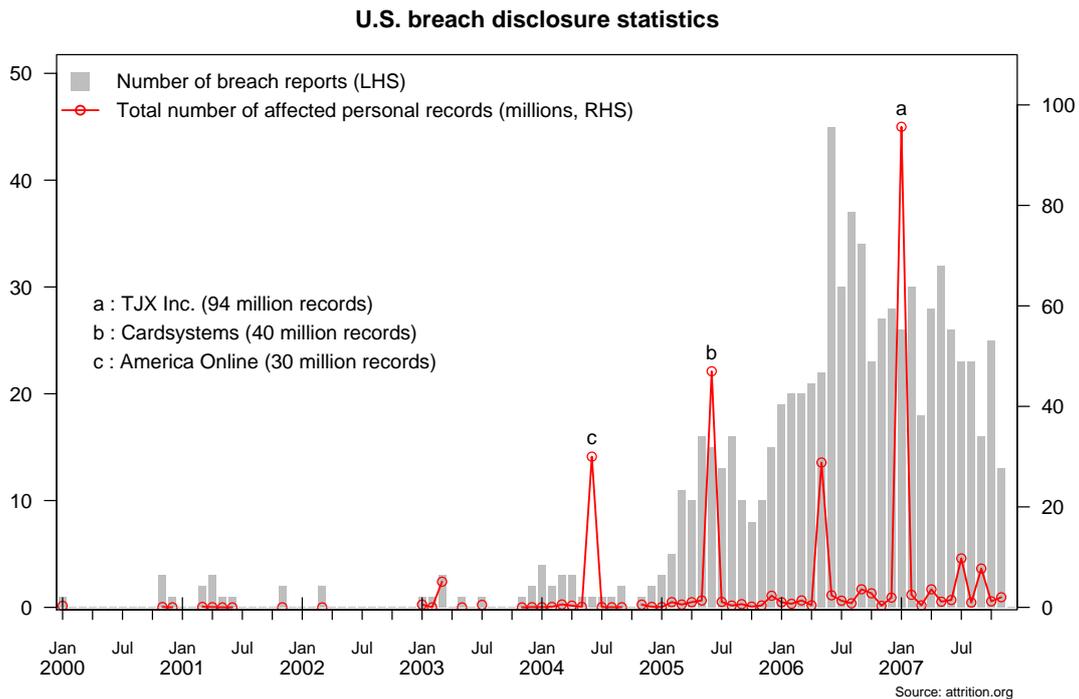


Figure 3: Bulge of breach reports after the introduction of disclosure laws in the US

gives brief details of each one [114]. The number of records compromised now exceeds 215 million. Several research groups, above all the contributors to attrition.org², a non-profit security resource page, are collecting the notifications that are sent, and it is to be expected that this data will provide a rich resource for future academic work understanding the nature of security breaches.

Figure 3 shows the monthly time series of reported breaches and affected personal records in the US since 2000. The rise from 2004 onwards demonstrates the breach notification legislation's impact. The distribution of the number of affected personal records has a very long right tail of a handful of landmark breaches with several million affected records. The exact shape of the left tail of the distribution may be distorted because many small breaches are silently mailed to the affected persons without attracting media attention. Only a few US states (e.g., New York) require breaches to be reported to a central data collection entity [103]. The median, as a robust measure, is a moderate 8,000 records per breach.³

Figure 4 shows the annual breakdown of breach disclosures by broad industrial sector (business, education, government and medical) and by breach type. Since 2003, the share of breaches due to hacking has continuously declined from more than 50% in 2003 down to 15% in 2007 (data up to and including November). Fraud and social engineering as

²<http://attrition.org>

³Note that these numbers are not additive across breaches as double-counting cannot be controlled for. Survey data published by Vontu, a vendor for data loss prevention solutions, suggests that the number of affected persons in the US is around 60% of the adult population for breaches up to mid-2007 (see <http://www.vontu.com/consumersurvey/>).

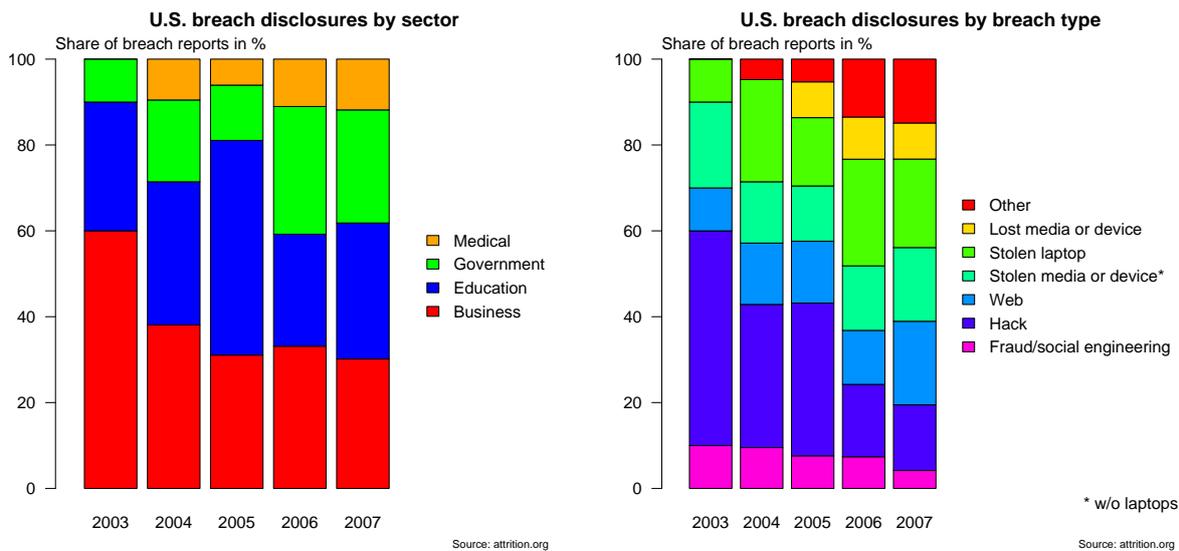


Figure 4: Distribution of breach reports across sectors (left) and breach types (right)

a reason for exposure of personal data is declining as well, although from a much lower level.

Figure 5 plots the distribution of breach types across sectors. Hacking is most prevalent in obtaining educational data whereas medical records are usually stolen. While these breakdowns were made on the basis of reported events, Figure 6 breaks down by the number of affected records. We computed the logarithm before calculating sector and type averages to account for the great variation in the number of records disclosed.⁴ Data losses are increasingly caused by accidents, despite the improved availability of full-disk encryption. Hacks account for a diminishing, but still substantial, portion of lost records. Notably, breaches via the web compromise the fewest records. As to the sectoral distribution, businesses tend to put most records at risk, while the education sector exposes the fewest. These plots demonstrate how data breaches can inform decision-makers of the biggest threats, along with their evolution over time.

In Europe, a security breach notification law has been put forward as a part of the 2007 review of the framework for electronic communications networks and services [46]. This would require notification to be made where a network security breach was responsible for the disclosure of personal data. This is a very narrow definition (necessarily so because it is being put forward specifically for one sector) and will only deal with a small fraction of the cases that a California-style law would cover.

The specific example we discussed above – of an automatic teller machine (ATM) being fitted by criminals with a skimmer that steals card details – would be covered by a California-style law. The bank would be required to notify every customer who’d used that machine during the period in which the skimmer could possibly have been in use, regardless of whether they were one of its customers or not. UK banks have resisted such

⁴A check for robustness using the sample median as aggregation function conveys essentially the same message; hence we omit the chart.

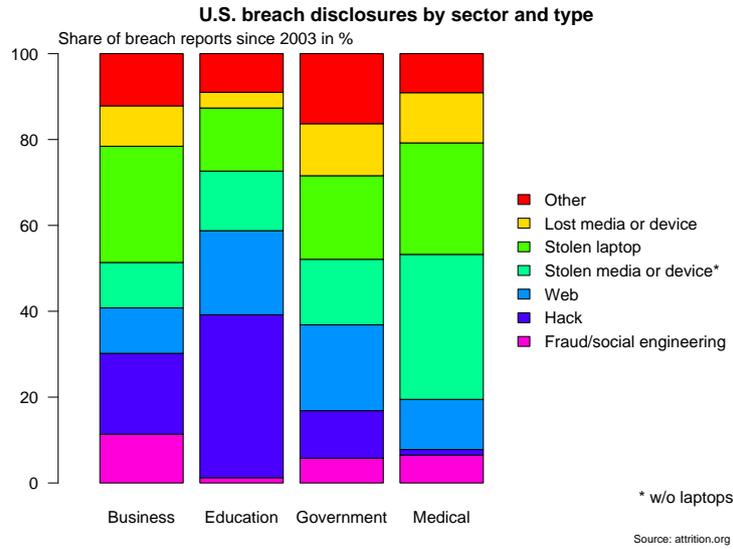


Figure 5: Breakdown by sector *and* breach type: Education is primarily hit by hacks while theft dominates in the medical sector

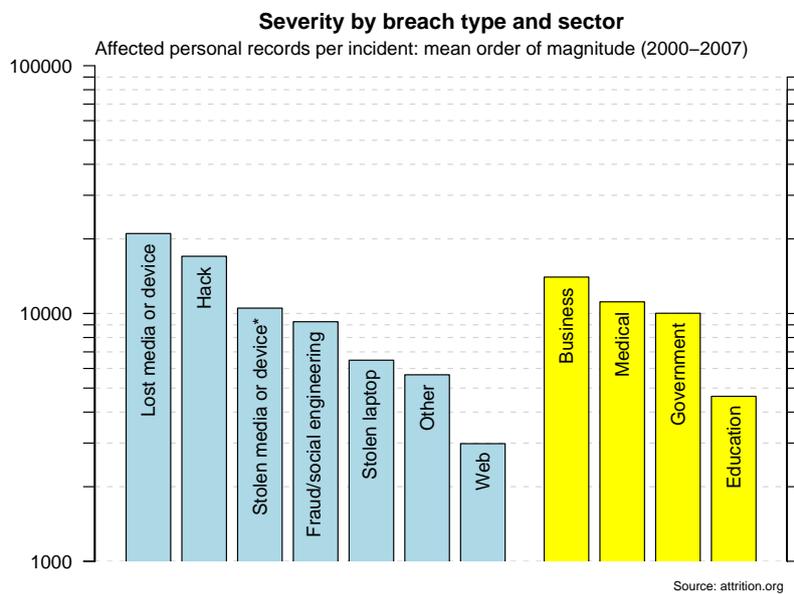


Figure 6: Log average of number of affected personal records per breach report broken down by breach type (left bars) and sector (right bars). Note the log scale.

proposals, claiming it would be inconvenient to contact other banks' customers. Yet their US operations appear to have no difficulty in complying with the law there.

In the UK, the House of Lords Personal Internet Security inquiry [76] recommended that the UK bring in a security breach notification law, and they made some recommendations on the detail as to how it should work. The report noted that there was no necessity to wait for a European Directive, but that steps could be taken immediately, and it was unimpressed by the telecom-only proposal from the Commission. The Government's response was negative, though they too didn't like the Commission's sector-specific proposal [133].

Although the House of Lords Select Committee saw advantages in bringing in country-specific legislation, the US experience demonstrates the disadvantages of a patchwork of local laws, and the obvious recommendation is that a security breach notification law should be brought forward at the EU level, covering all sectors of economic activity. The current EU proposal applies only to telecomms companies, and so would not solve the ATM problem, or for that matter the ChoicePoint problem – unless the ATM operator, or data aggregator, were owned by a phone company. There was a solid consensus among stakeholders that the law should not discriminate between economic sectors.

The point of security breach notification is to avoid all the complexity of setting out in detail how data should be protected; instead it provides incentives for that protection. Thus it does not impose the burden of a strict liability regime across the whole economy (though in many sectors this might be desirable), but relies on 'naming and shaming' to provide encouragement to firms to improve the protection of personal data. Competent firms have nothing to fear from breach notification, and should welcome a situation where incompetent firms who cut corners to save money will be exposed, incur costs, and lose customers. This levels up the playing field and prevents the competent being penalised for taking protection seriously.

Recommendation 1: We recommend that the EU introduce a comprehensive security-breach notification law.

It is important that the law be as effective an incentive as possible, and lessons can be learnt from the US regarding this. As well as informing the data subjects of a data breach, a central clearing house should be informed as well. This ensures that even the smallest of breaches can be located by the press, by investors, by researchers, and by sector-specific regulators. The law should set out minimum standards of clarity for notifications – in the US some companies have hidden the notifications within screeds of irrelevant marketing information. Finally, notifications should include clear advice on what individuals should do to mitigate the risks they run as a result of the disclosure; in the US many notifications have just puzzled their recipients rather than giving them helpful advice.

4.2 Metrics

There has for many years been a general lack of adequate statistics on information security. The available data are insufficient, fragmented, incomparable and lacking a European perspective [69]. Depending on the source and mode of data collection, further issues emerge, such as intentional under- and over-reporting as well as all kinds of unintentional response effects. Vendors in particular have often played up the threats, for example by

claiming that banks and other firms report only a small fraction of incidents in order to avoid losing public confidence. Indeed, one of us (Anderson) recalls working for a major bank in the late 1980s and truthfully assuring security product salesmen that only a small number of modest losses were sustained as a result of electronic crime – and being accused of lying.

Crime statistics are a notoriously hard problem even in the non-electronic world. Governments and police forces have every incentive to find ways to discourage the reporting of minor crimes and to change procedures to minimise numbers. As a result, one gold standard is the victim survey, whereby a sample of members of the public are asked every year whether they have been a victim of crime and if so, what. Electronic crime is no different. In the next section, we present a principled analysis of what statistics should be collected, why, from whom, and how.

4.2.1 What are the statistics for?

The primary value of statistical data, and the main justification for its collection by government agencies, is to mitigate information asymmetries by generating useful signals for economic decision making, whether by policymakers, firms or individuals. The signal-to-noise ratio of a piece of information may vary between different stakeholders; each purpose also sets specific requirements for data accuracy, frequency and timeliness.

Individuals and organisations benefit from data on security properties when making consumption and investment decisions. For this purpose, it is important that the data is disseminated in a timely manner, reflecting the fast pace of technological development. Indicators must also be broken down to meaningful categories, e.g. by suppliers or product lines. For example, it is not overly helpful to this group to know how the total number of vulnerabilities found over the past couple of years, but rather how Windows compares to Mac OS X or Linux, or how Internet Explorer compares to Firefox. (Such data are already available but are easily accessible only by security and IT professionals.)

Security professionals at organisations and infrastructure providers (e.g. ISPs) need statistical data to plan and implement appropriate protection and to react to current levels of threat. Again, timeliness, accuracy and breakdowns by product are important. Trends also matter, while indicators should be updated as needed to cover new and unknown events. (One user survey of security failures is less useful than it could be because it lumped together the virus and worm infections common in the early 2000's with the phishing attacks that have grown rapidly since 2004.)

Another key use of statistical data is *policy formation*. Completeness, consistency and comparability both across reference areas and time are more important for this than timeliness or level of detail. Good statistics can support policy evolution in a number of ways; for example, the most effective policies of individual Member States can be identified and become best practice (as security-breach disclosure laws have spread across the USA). This process is prone to be very slow, as governments (rightly) do not act on the short technology cycles of the IT industry. Testing different policies in a 'natural experiment' is likely to provide better outcomes than too-early harmonisation.

Consistent, comparable metrics enable greater *transparency*. At present, there is great variation between organisations in their security practices. For example, Moore and Clayton have studied the effectiveness of phishing website removal countermeasures instigated by the banks and specialist 'take-down' companies [101]. They have found

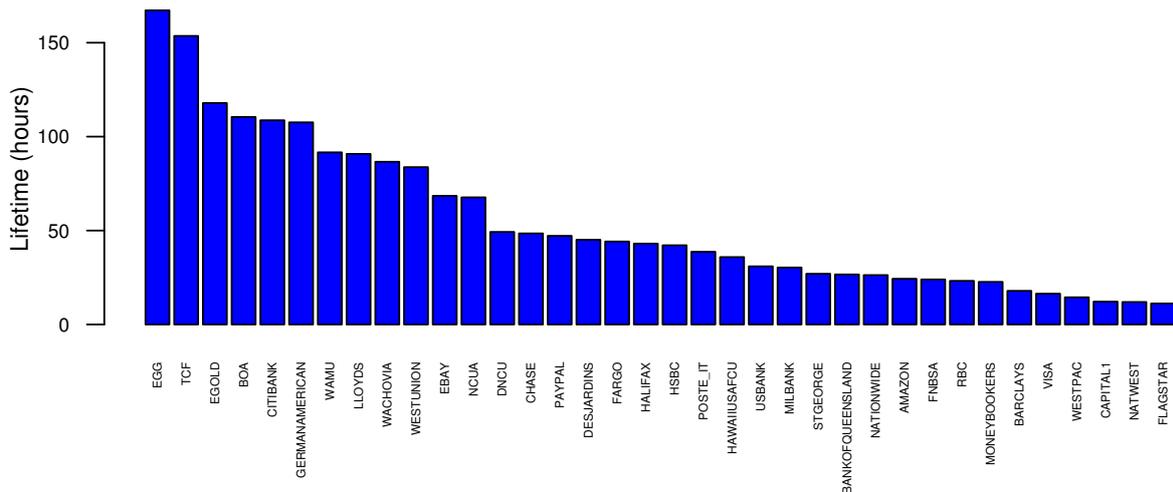


Figure 7: Phishing-site lifetimes per bank. Source: [101]

that the performance of banks and the responsiveness of ISPs is very skewed, with the best outperforming the worst by more than one order of magnitude. Figure 7 shows the average lifetime of fraudulent phishing sites for each bank impersonated. This variation demonstrates a need for more comparable measurement across ISPs and banks – the laggards are weakening Internet security, and they get away with it because there is no transparency to hold them to account!

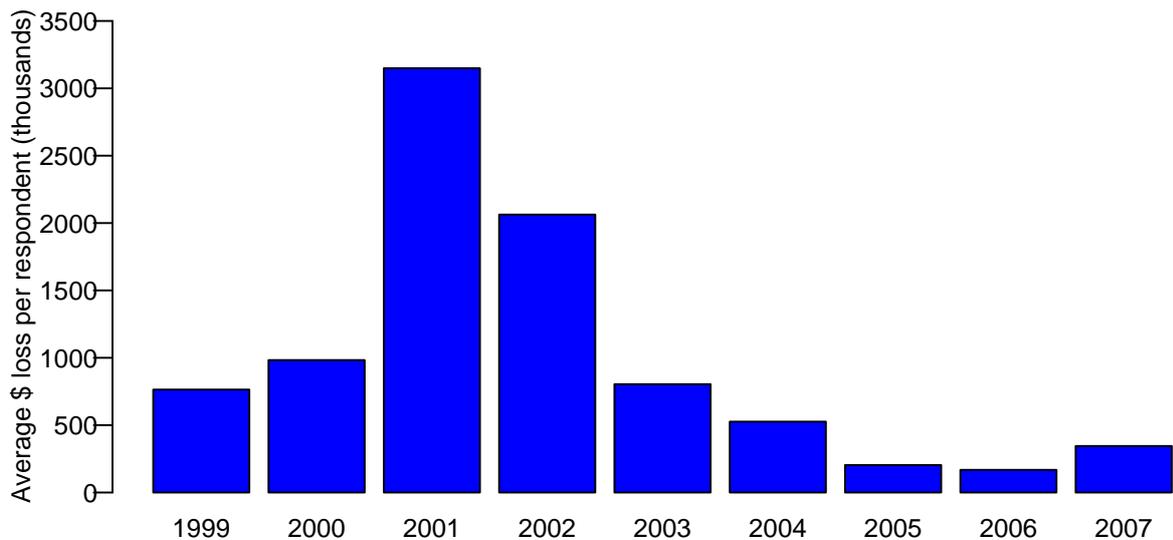
Academic research is another consumer of statistical data. The recent rapid growth of security economics has explained many failures that were previously puzzling [6]; better data will bring further rewards in the form of deeper understanding that will in turn lead to better policies in the long run. While aggregated data allows comparisons between countries, the statistical power of such studies is rather weak and prone to third variable problems (the ‘ecological fallacy’ [116]). Hence, micro-data should be made available for research. Timeliness is less important, but higher frequency time series help to improve the reliability of inference.

4.2.2 What statistics are already being collected?

Many organisations already collect and analyse statistical data on Internet security. In fact, ENISA has recently published a report that outlines over 100 sources of data on information security [21]. Published data comes in many forms.

One common approach is to conduct *surveys*. For the past twelve years, the US-based Computer Security Institute has annually surveyed enterprises, asking respondents whether they have been attacked and, if so, what the resulting losses were [28]. We examine data from the CSI survey in Section 4.2.3. There has so far been one Community initiative at collecting statistical data relevant to this report. In 2003, Eurostat started collecting data on Internet security issues from both individuals and enterprises in its ‘Community Surveys on ICT Usage’ [45]. Data from the Eurostat study is presented in Appendix A and discussed in Section 4.2.3.

Security breach-disclosure reports provide another useful data source. Groups such as



Source: CSI 2007 Computer Crime and Security Survey

Figure 8: Average annual reported losses per enterprise attributed to computer crime.

`attrition.org` collate reports, which we discuss in Section 4.1 above.

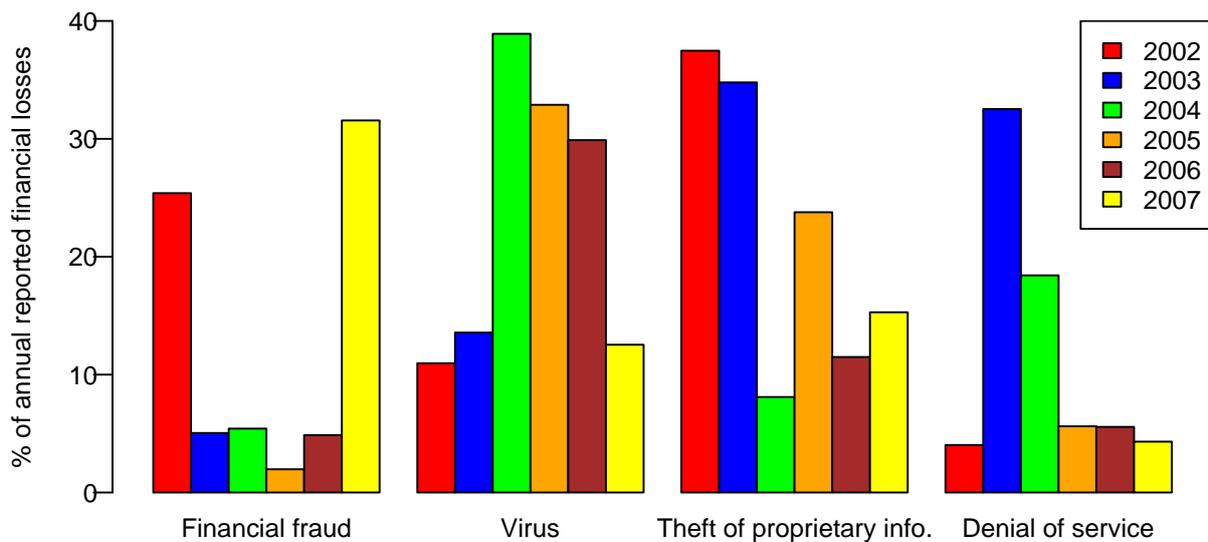
Direct observation is a third category of data collection on information security. Many security vendors regularly publish reports. For example, Symantec has published a semi-annual Internet Security Threat Report (ISTR) since 2002 [130]. Symantec directly measures many types of malicious activity using its global infrastructure of 40,000 sensors. We discuss Symantec’s report in Sections 4.2.3 and 6.5. Other security organisations such as McAfee [94], SANS [118], IBM [81] and Microsoft [99] have also published useful reports on attack trends. Industry groups also sometimes disclose useful statistics, including the Anti-Phishing Working Group (see discussion in Sections 4.2.4 and 4.3.2) and APACS, the UK payments association (see discussion in Section 4.3.2). Finally, some academics conduct useful data collection and analysis. In this report, we refer to analysis of phishing websites by Moore and Clayton [101] in Figure 7 and malware tracking by Zhuge et al. [138] in Sections 6.5 and 8.2.

4.2.3 Case studies of security statistics

In this section, we discuss just three promising, regularly published reports, which might serve as a useful data source: the Eurostat survey, CSI survey and Symantec report. It is hoped that by studying these examples in greater detail we can demonstrate both the opportunities and challenges presented by existing data collection efforts.

What each data source does well The CSI survey⁵ has done a good job of asking questions consistently over long time periods. The survey has added and removed a few questions from the report over time, but many of the fundamental questions remain unchanged. It also is unique in that it asks respondents to report their estimated monetary

⁵This was called the CSI/FBI Computer Crime and Security Survey until 2007.



Source: CSI 2002–2007 Computer Crime and Security Surveys

Figure 9: Proportion of annual reported losses attributed to different threat categories.

losses due to various attack types. While there are undoubtedly problems associated with asking firms to self-report losses, there is no better measure of monetary losses at present.

Figure 8 plots the reported annual average loss per responding firm. Notably, the average loss shot to a peak in 2001, while decreasing substantially in subsequent years. While the exact figures may not be generally applicable, the downward trend in individual firm losses may be.

To further demonstrate the benefits of good time-sequenced data, we have collated the financial loss figures broken down according to type from the CSI survey. Figure 9 plots the percentage of reported annual losses for four types of attack, collated from six years of CSI surveys. This figure demonstrates how the biggest security threats to firms can quickly change from year to year. In 2002 and 2003, the worst losses were caused by theft of proprietary information. However, in 2004 the losses attributed to viruses shot up to become the largest cause, at nearly 40% of losses. Viruses continued to be attributed as causing the most losses in 2005 and 2006. Meanwhile, the losses due stealing proprietary information fell sharply. Finally, in 2007 financial fraud, which had accounted for less than 5% of all losses in 2003–2006, accounted for around a third of all losses, displacing viruses as the loss-leader. What this data shows is that the cause of losses is difficult to predict based on what caused them in previous years. In fact, the strongest conclusion one can draw from the graph is that the cause of losses is likely to continue changing rapidly.

The Eurostat survey is beneficial because it surveys consumers as well as enterprises, and also provides comparative data between the responses in different European countries. Annual data are available for a broad (but still incomplete) set of Member States on the percentage of individuals and enterprises with Internet access who have:

- encountered security problems,
- taken ‘ICT security precautions’ within the last three months,

- installed security devices on their PCs and updated them within the last three months.

Individuals are further broken down by age group and residence (urban vs. rural areas). Data on enterprises is available for different firm sizes and NACE main sectors (excluding the financial sector). A subset of the data appears in appendix A of this report.

The Symantec report is based on direct measurement of malicious Internet activity. The advantage of this approach over a survey-based one is that it overcomes problems of respondents' differing understanding of what threats are. They have also appreciated the value of collecting data in a consistent manner over time. Unlike many other vendors that publish data, they make all past reports publicly available, and they normally describe methodological differences between previous reports when necessary. They also make a substantial effort to describe their methodology in appendices.

Problems with the data sources Unfortunately, none of these existing data sources is without problems. The CSI survey is better at asking questions consistently and reporting in the same manner; Symantec's ISTR is worse at this. Even when the measured statistic remains unchanged, the presentation may change dramatically from report to report.

While the CSI survey has done well to produce loss figures, there are major issues with loss assessment. In some jurisdictions, police will not pass a crime to a specialist unit (such as a computer crime squad) unless the losses pass some threshold. So a company that has been the victim of a hacking attack will seek to maximise its apparent losses, for example by claiming that the disruption caused by the clean-up must have cost an hour's productivity from every staff member, and then multiplying this by their charge-out rate. Had they been making an insurance claim, the loss adjustor might only have allowed the extra overtime worked by system administrators. The gap between the two figures can be more than one order of magnitude.

Response effects can include addressing the survey requests to 'the computer security manager' or 'the chief internal auditor' with the result that responses are obtained only from firms large enough to have someone in that role, or sufficiently interested to at least read the letter. The majority of respondents to the CSI survey, for example, have over 1,500 employees and turnover in excess of USD 100 million. And the recent rapid growth in attacks on individuals, rather than companies, increases the effective bias of surveying large-company officials.

One problem with Symantec is that it is not an unbiased reporter. Their reports are published principally for marketing reasons. It is not surprising, then, that sometimes their data are consistently over-reported and later revised down, which gives the false appearance that a particular problem is rising.

To demonstrate over-reporting, we studied more closely one statistic measured in several reports, which tracks the proportion of malicious code that exploits confidential information. In volume 12 of the report, covering January to June 2007, 65 % of malicious code exploits confidential information, compared to just 53 % in the previous six months. However, the earlier report claimed 66 % for this period of July to December 2006. The potential for such discrepancies are mentioned in the report's appendix: 'there may be slight variance in the presentation of the same data set from one volume of the *Internet Security Threat Report* to the next'.

Table 1: Proportion of top threats that undermine confidentiality according to different Symantec reports. The most recent four reports overstate the proportion, which repeatedly gives the false impression that threats to confidentiality are rising.

Report period	Reference period								
	2007		2006		2005		2004		2003
	H 1	H 2	H 1	H 2	H 1	H 2	H 1	H 2	
H 1 2007	65 %	53 %							
H 2 2006		66 %	48 %	55 %					
H 1 2006			60 %	60 %	40 %				
H 2 2005				80 %	74 %	54 %			
H 1 2005					74 %	54 %	44 %		
H 2 2004						54 %	44 %	36 %	
Revised	– ¹⁾	53 %	48 %	55 %	40 %	54 %	44 %	36 %	

Note: H 1 and H 2 denote halves of the year. ¹⁾ pending

Source: Symantec

We examined earlier reports and indeed found more discrepancies. Each row in Table 1 gives the proportion indicated by each volume of the report, while each column indicates the time period. Notably, every report suggests that malicious code is increasingly attempting to expose confidential information. However, each of the four most recent reports revised down figures for earlier periods. Accounting for these revisions, the proportion no longer appears to be increasing, but rather vacillating around 50%. Whether intentional or coincidental, more care is required to draw meaningful conclusions from vendor-provided data analysis.

Another problem is that the conclusions from the different reports often disagree. For example, according to the CSI surveys (see Figure 9), losses attributed to denial-of-service attacks peaked in 2003 at one third of all losses, fell sharply to 18% in 2004 and under 6% of all losses in 2005. The Symantec ISTR, by contrast, paints the opposite picture in a graph of observed denial-of-service attacks in 2004 and 2005 (Figure 10). They observed very few attacks in 2004 (less than 100–200 per day), which increased massively in 2005 (up to 1,600 per day). Despite this increase in attacks, the losses attributed to them in the CSI survey fell dramatically. This could be due to poor sampling by the CSI survey, or it could be that the increase in attacks had no bearing on the damage inflicted. Regardless, it demonstrates that merely counting attacks without assessing the associated costs can be misleading.

The Eurostat survey has also been plagued with difficulties. While the definition of indicators appears very reasonable, collecting reliable responses in surveys has turned out to be problematic. Most items were discontinued in 2005 as the current questions are inadequate.

Especially for the household survey, respondents are not technically proficient enough to comprehend the questions or know whether their computers are protected and updated [45]. This potential measurement error becomes evident in Table 9, where for some countries the percentage of households who claim that they have updated their virus

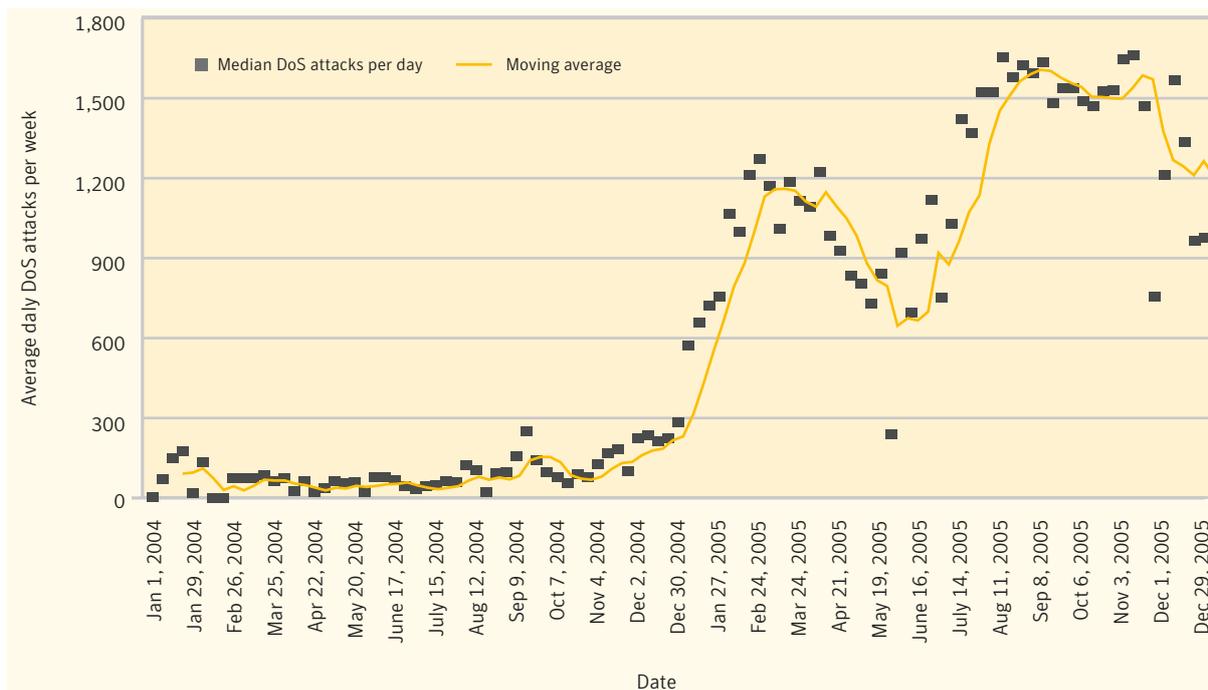


Figure 10: Denial-of-service attacks per week in 2004 and 2005 (Reproduced from Symantec Internet Security Threat Report IX)

checking program in the last three months exceeds the percentage of household who report that they have installed one. A closer look at the tables reveals further potential data issues. Remarkably, the only indicator that has ‘proved feasible’, according to an EC document on benchmarking [45], contains hardly any variance and consequently conveys almost no information.

Nearly all enterprises reported that they have taken precautions (see table 10 in the appendix; percentages of enterprises ‘having taken precautions’). It remains unclear if the term ‘precautions’ was simply too broad and vague, so that every respondent could actually find at least one measure of precaution, or if the high values result from (rational) over-reporting or (behavioral) acquiescence effects [95]. Consequently, there is room for more research on constructing valid indicators on information security for both household and enterprise surveys.

Against this backdrop, it is unfortunate that Eurostat’s work on these indicators appears stalled, with the dim outlook of a special module on security, scheduled for 2010 – at the last possible date of the current i2010 agenda [45]. At least for the enterprise survey, Eurostat should be able to achieve a similar degree of detail as the CSI survey.

4.2.4 How should statistics be collected?

Having just discussed the ‘state-of-the-art’ in statistical indicators for information security, we next describe how statistical indicators should be collected. We also note additional challenges to developing good statistical indicators.

Required level of detail and breakdowns Metrics quantifying properties of *security providers* should be broken down to the individual provider, product or service. This

enables security consumers to make informed investment decisions. Additional aggregates across different indicators can reduce noise, facilitate decision making, and add value to signals. Implementing best practice for the choice of weights and in particular the aggregation function should ensure that composite indicators are calculated in a fair and hard-to-manipulate manner [105].

Metrics quantifying the behaviour of *security consumers* should be available on a higher level of aggregation, e.g., by region and consumer type. This enables cross-sectional and longitudinal analyses for comparing policies and attitudes while preserving the individual respondent’s anonymity. This is essential to discourage opportunistic responses and to paint an objective picture for policy makers.

Quantitative metrics of individual *security incidents* related to third parties’ data (e.g. customer data), which may be collected through complimentary breach disclosure legislation (cf. Section 4.1), should be released as detailed as possible to enable concerned consumers to take precautions in the event of data compromise and to encourage data brokers to take precautions.

Table 2 provides a summary of our recommendations for breakdowns and supplemental information.

In specific applications, there may be debate over collecting further data. In the bank fraud data example, researchers (and policymakers) might like to know whether a particular bank’s systems – from its technical systems to its dispute-resolution procedures – discriminated against less educated citizens, and so there may be a strong case for a breakdown of dispute data by both provider and consumer type. Robustness to regulatory differences across Member States can be achieved by measuring the proportion of customer claims that are refunded or rejected. Robustness to social bias can be achieved by measuring how the refund ratio varies with social indicators across Member States.

Challenges to constructing statistical indicators

Challenge 1: Definition of robust indicators. Defining meaningful metrics for information security is particularly difficult because of the dynamic conflict between attack and defence. Attackers adapt very quickly, so metrics defined for a particular tactic may lose relevance over time. For example, a metric for online identity theft defined as the number of victims deceived by phishing sites stops working as a signal as attackers move from phishing towards installing malware equipped with key-loggers (see Figure 11). This example also demonstrates the need for high-frequency time-series data; here it is needed not just for academic research after the fact, but also as an operational matter for crime-fighters and service providers.

Robustness to short-lived tactics can be achieved by measuring losses due to cyber-attacks. For the case of losses to phishing and keyloggers, a better measure is the number of customer disputes in online payments, or perhaps the total disputed transactions in euros. This illustrates the value of combining data from multiple sources – communications service providers, security vendors and financial institutions.

Challenge 2: Definition of reference objects. A side-issue of compiling product-specific security indicators is the definition of eligible products and services. Market size could serve as a criterion, but it brings the difficulty to measure market size objectively,

Table 2: Recommended breakdowns for information security indicators

Data subject	security providers		security consumers		security incidents
	products	services	enterprises	households	
Data fields for breakdowns					
geographical region		✓	✓	✓	(✓)
product / service name	✓		×		(✓)
version	✓				
module ^{a)}	(✓)				
platform	✓	✓	✓	✓	(✓)
industry ...					
NACE sector		(✓)	✓		✓
IT dependence			✓		✓
NIS sensitivity			✓		✓
firm size			✓		✓
security spending ^{b)}	(✓)	(✓)	✓	(✓)	(✓)
losses due to bad security ...					
data			✓	(✓)	
time			✓	(✓)	
money			✓		✓
socio-economic indicators					
IT literacy				✓	
attack vector					✓
type of data affected					✓
Supplemental information					
# data records affected					✓
# individuals affected					✓
contact details	✓	✓	×	×	✓

Legend: ✓ = recommended, (✓) = optional, × = don't disclose, # = number of

^{a)} for composite products, e.g. to distinguish the firewall from hard disk encryption of an OS

^{b)} R&D for providers, investment for consumer enterprises, yes/no for households

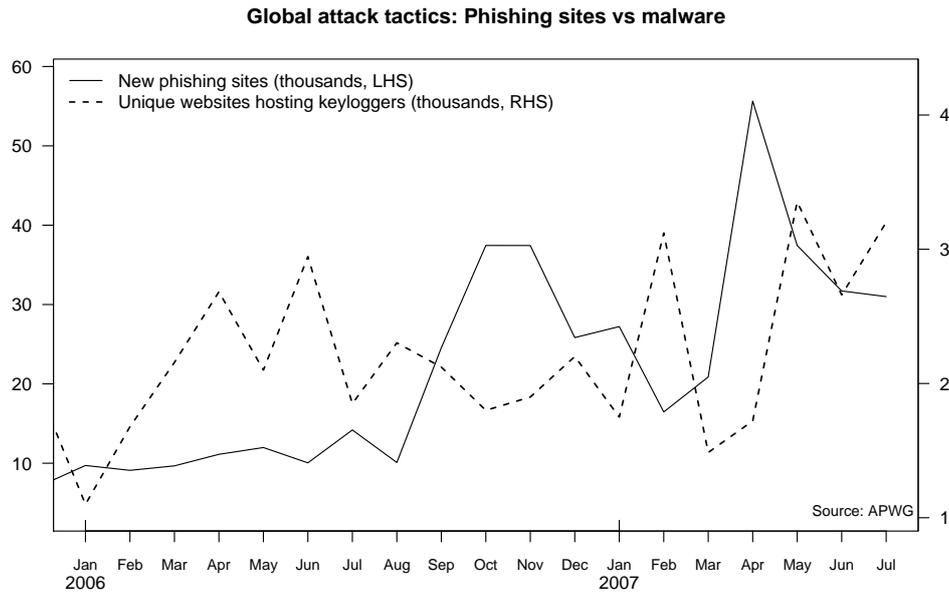


Figure 11: Attack trends vary rapidly over time. Phishing sites and keyloggers are substitutes, the Pearson correlation of first difference series is -0.4 .

and it may discriminate against products or services at the beginning of their life cycle. This can turn into an additional market entry barrier for newcomers, thus discouraging innovation. Even worse, incumbent firms can easily reach the threshold market size for their new products by bundling them with established products for a limited period of time. We are not aware of a perfect solution, so data-collecting authorities should at least be made aware of the issue. A reasonable compromise could be to monitor a larger set of products at a lower frequency (e.g. annual reports or reviews) while concentrating on a narrow set of main market players in more frequent and timely regular updates.

Challenge 3: Identification of reference areas. Another possible fallacy lies in the assignment of reference areas to traffic data. Although most IP addresses can be assigned to individual countries quite reliably, an offending source host in a particular country does not mean that an attack originates from this country, since the true perpetrators may hide behind a cascade of proxies in various countries to conceal their traces.

Challenge 4: Seasonality and calendar effects. Some methodological challenges also remain on the processing stage. Seasonal adjustment, while being standard for almost all economic time series, is hard to apply to observed traffic data, as seasonal patterns are superimposed with short-lived or transitional market trends (e.g. shift from YouTube traffic to Skype traffic), or tactics in attacker behaviour. What's more, seasonal factors in attacker behaviour appear to be calendar-driven, as anecdotal evidence of 'school holiday effects' in website defacement suggests. Contrary, seasonality in attacks against unpatched systems is determined by the patch release pattern of software vendors (e.g. 'Exploit Wednesday' following Microsoft's 'Patch Tuesday'). Other data sources exhibit more

regular patterns, which can be dealt with using conventional adjustment methods. For example, spam advertising phishing websites falls at the weekend [112], while the sites themselves, by contrast, are launched with no such discernible pattern [101].

4.2.5 Metrics derived from market price information

Market prices are formed in negotiation between agents, who adjust their behaviour based on expectations about the future. Market prices thus contain some information about agents' expectations and may serve as forward-looking indicators – in contrast to the above statistics, which are calculated from information about historical states of the world.

The efficient market hypothesis, in its strong form, suggests that stock market prices aggregate all information relevant for forming expectations of future profits [64]. If information security matters for a company, then the stock market valuation should react to news about security incidents. Several authors have conducted event studies, the method of choice for analyzing market reactions to news. They found measurable negative market price reactions following reports of denial-of-service attacks against popular websites [63, 77], security incidents [20, 71], computer virus contagion [78], vulnerability disclosure in software products (though smaller in magnitude) [131], and privacy breaches [2]. Another study also finds positive market reactions for listed security firms when news on security breaches is in the media [23]. Other co-variates have been examined, such as firm size (smaller firms suffer relatively more), business model (B2C firms suffer more from denial-of-service than B2B), and exposure of confidential data (which amplifies the market reaction).

Aside from the potential for publication bias, event studies offer evidence that information security news impacts markets. However, the general usefulness of stock market prices as a direct metric for security properties is limited for two reasons: first, event studies capture short-term losses only and it is hard to estimate medium and long-run impact of security failures on stock prices, and second, opportunities to conduct event studies are limited to the occurrence of extreme events, such as attacks and security breaches. In general, stock prices aggregate too much diverse information to be useful as a security indicator. As a consequence, researchers have studied markets closer to the object of interest, to isolate information security signals from the noise in general price information.

A recent proposal [67] to track 'underground market indices' in IRC channels to feed forecasting and threat prediction tools may sound a bit daring: these markets operate as exchange platforms for stolen credit card and identity details, hacked accounts, spam distribution and related services. Yet the 'street price of drugs' is commonly used as a signal for the effectiveness of enforcement; and in information security, too, the price of contraband goods such as stolen credit cards is considered important by some players. However, the signals are not straightforward. Officials of one bank remarked to one of us that a fall in the 'street value' of their credit card numbers to under a euro was a good thing. They believed that this was not signalling that the market was flooded with their customers' credentials, but rather that their back-end fraud-prevention mechanisms were good enough to prevent significant value extraction from a stolen credit card number alone.

Many ideas have been put forward to use markets to extract security-related information – most of them having been designed to counter security market failures. They seek not just to align incentives, but also to provide new security metrics.

The literature distinguishes various forms of so-called ‘vulnerability markets’, of which various kinds can already be observed in practice [15, 129, 100]: black markets, vulnerability brokers, bug bounties, and bug auctions. There are also suggested innovations, such as exploit derivatives. Meanwhile, the price of cyber-insurance provides an indirect market measure of overall systems vulnerability. We will now look at each of these briefly.

The vulnerability black market is a catch-all term for the unregulated vulnerability markets, which are how some security researchers currently extract revenue from discovering flaws. Although referred to as a ‘black market’, the business per se is not illegal under most jurisdictions, although selling an exploit to someone the researcher knows is likely to make criminal use of it is an offence in most countries, as is blackmailing a vendor. Selling an exploit to a national intelligence agency, or to a firm that sells keylogger software to police forces, or to a firm that reverse-engineers protected software on behalf of lawyers conducting intellectual property disputes, or even to a firewall vendor who wants advance warning of exploits, is generally legal.

However, market participants point to a lack of transparency in pricing; difficulties in finding buyers and sellers; and possible difficulties faced by a seller in ensuring a buyer’s bona fides. Successful dealmaking largely depends on personal contacts. The market also suffers from the typical problems of information goods. Vulnerabilities are often easily-duplicated experience goods, in that a seller who demonstrates one to a potential buyer may give away the secret. Also, a recipient may sell a vulnerability onward, and a seller may sell the same vulnerability to multiple buyers despite giving each of them an assurance to the contrary. (Such contracts are difficult to enforce as vulnerabilities are often independently rediscovered.) Thus ownership is hard to establish, and exclusivity is hard to guarantee. Transactions in these markets are further impaired by the rediscovery risk making vulnerability information time-sensitive [111]. All in all, market participants describe the current state as sub-optimal [100]. Finally, these illiquid markets not only have high transaction costs, but prices are rarely publicised, which greatly limits their usefulness as security metrics.

Vulnerability brokers act as infomediaries in the ‘black market’. Four players – iDefense, TippingPoint, Digital Armaments and Netragard, all based in the US – have offered to pay security researchers a lump sum for vulnerability information and distribute it among a closed group of subscribers to their alert services, which are sometimes bundled with filtering or intrusion detection (IDS) tools. Despite some competition, the prices that these brokers pay for vulnerabilities are a magnitude smaller than on the black market,⁶ so security professionals are concerned that the big discoveries are still being sold to criminals [129]. A welfare-economic analysis concludes that vulnerability brokers are not socially optimal, since users who do not participate in the closed circle of subscribers cannot protect their systems in time [85]. In addition, subscribers might leak pre-patch

⁶Prices on the black market are reported to range between 5 and 6 digit dollar amounts, whereas brokers typically pay 4 digit compensation plus bonuses for frequent contributors. According to industry sources, the typical contributor is a freelance security researcher based in central or eastern Europe.

vulnerability information to miscreants. Finally, the brokers do not disclose the prices they pay for vulnerabilities; so this market structure is unsuitable for deriving security metrics.

Bug bounties are the oldest form of vulnerability market. The idea is that software vendors offer a cash reward of a preset value for bug reports, to stimulate researchers to look for and report bugs, and also to boost public confidence in the vendor's product. The most famous example is Don Knuth's quadratically increasing (but capped) reward for reports of errors in his \TeX and METAFONT software [89]; another is the Mozilla Foundation's fixed reward of USD 500 for critical security bugs in the Firefox browser [102]. The reward can be one measure of security strength [119], although Knuth reckons that most of the people to whom he's sent a reward cheque have never cashed it – they leave it in the wall as a trophy. In effect, the problems of bug bounties are indicator quality and vendor commitment. Recently, some of the vulnerability brokers have started offering bounties where the vendors don't: TippingPoint offered a USD 10,000 reward for a zero-day exploit on a MacBook Pro, which attracted a researcher who found a bug in QuickTime [86].

Bug auctions were proposed as an extension to the price-setting mechanism of bug bounties. The idea is to draw on auction theory to formalise and improve the mechanism design [110]. They can be either ad-hoc or formal. Researchers sometimes offer newly discovered vulnerabilities on public auction websites⁷, but such auctions are often shut down by the site operators before termination. Such 'buyer-administered' auctions can generate useful price information if run frequently enough [15]. This opened an opportunity for a Switzerland-based company, WabiSabiLabi⁸ to set up and run a dedicated auction platform for vulnerability information since early 2007. In the period up to the end of 2007, its market history lists 32 successful sales at an average price of EUR 1840.

The above options suffer from the usual problems of dealing with information goods. We will now present two concepts for indirect vulnerability markets, which do not depend on passing sensitive information from a seller to a single buyer.

Exploit derivatives transfer the idea of binary stock options to vulnerability discoveries [15]. Exploit derivatives are contracts with a defined par value, date of maturity, software and platform specification. The par value is payable to the owner on maturity if there exists an exploit against the defined product and configuration. The existence of an exploit can be attested in a verifiable way. As in other prediction markets [136], exploit derivatives can be issued in pairs: a 'vulnerability contract' that pays if the exploit exists, and a 'security contract' that pays otherwise. Owners can trade contracts on an exchange at public prices – which can be interpreted as the market's best estimate of the probability that a security hole will be found. Exploit derivatives may convey two other useful properties. First, a good researcher's information advantage could yield much higher profits than most of the rewards seen so far in existing market types. Second, exploit derivatives can be used to hedge the risk of being exposed to a particular software vulnerability. This opens the door for other parties than security researchers and software vendors to become legitimate and welcome market participants, thus boosting liquidity.

⁷see for example <http://it.slashdot.org/article.pl?sid=05/12/12/1215220>

⁸<http://wslabi.com>

Cyber-insurance contracts transfer risk from firms to insurance companies that agree to cover any financial loss incurred through damage to or unavailability of assets caused by computer security incidents. Insurers, over time, accumulate data on loss amounts and the effectiveness of safeguards which they mine to improve risk assessment and suggest best practice mitigation strategies to their clients [14]. In a competitive market for cyber-insurance, premiums would be affected by the insured’s level of security, both technical and managerial. This price information could serve as a security indicator just as car insurance premiums signal car safety properties (and typical driver behaviour). However, cyber-insurance markets suffer from a number of limitations, chiefly liquidity, competitiveness and price transparency. Some firms have some cover in their general insurance policies, for example against negligence or dishonesty by employees, and some buy special cover. (We will discuss cyber-insurance in much greater detail in Section 9.1.)

To sum up, there is a rich diversity of concepts for vulnerability markets, which makes it hard to predict which models will thrive. Nevertheless, a theoretical comparison of the market models reveals that some are better suited than others to overcome information asymmetries, align incentives, and balance risk. Notably, cyber-insurance and complementary financial instruments for risk sharing (see also Section 9.1) are socially beneficial and thus clearly preferable to rather obscure and unregulated black markets. Information security legislation, therefore, should not inhibit the proliferation of legitimate vulnerability markets. For instance, a recent feasibility study on exploit derivatives found plenty of legal obstacles in the US legislation (we are not aware of a similar study for the EU) [121].

4.3 Information sharing

While the primary aim of breach disclosure legislation is to encourage firms to adopt well-known security practices, this is not its only benefit. The type and frequency of attacks can inform other firms of the evolution of threat types, and thus help firms prepare defences before they are targeted. Breach data also helps all firms to develop techniques for detecting and preventing attacks. There are many types of attack, and security engineers have to learn from other’s failures as well as from their own.

Both qualitative and quantitative data may be shared between organisations. Some data are shared with the public, whether through news, technical alerts, or the research literature. In other cases, security information is shared in confidence between firms in the same industry. When quantitative data is shared across industries, it’s important to develop comparable metrics and to bear in mind the uses to which the data will be put (see Section 4.2).

4.3.1 Costs and benefits of sharing

The costs associated with sharing must also be taken into account. First, companies do not like publicising security breaches, because they might be exploited by competitors, receive a bad press, or get the company sued. Managers may also be loth to disclose breaches in case they get fired. A breach-disclosure law, which we advocate, will blunt these disincentives.

A list of reasons for not sharing data, derived from [70], is compiled in Table 3. Data from the 2007 CSI Computer Crime and Security Survey [28] report that negative publicity

Table 3: Barriers to sharing security information

- Loss of reputation and trust
- Risk of liability and indemnification claims
- Negative effects on financial markets
- Signal of weakness to adversaries
- Job security and individual career goals

is the most-cited reason to abstain from reporting a security problem (26 % of respondents name it).

Another worry about information sharing is that firms might free-ride off the security expenditures of other firms by only ‘consuming’ shared security information and never providing any [73]. For example, sharing information about an exploit in a commonly-used application that was discovered during penetration testing lets other firms improve their security without incurring the same cost of discovery. Even the threat of such free-riding can stymie sharing. Where there has been limited sharing, it may be down to these costs. Sharing sensitive security information could also, in some circumstances, provide a competitive advantage to firms receiving the information, for example by disclosing that a firm was working with some particular platform. However there is little evidence that this is a major concern in practice.

Firms sometimes object to sharing data for fear of violating privacy regulations. Some European ISPs claim they cannot look at individual IP addresses when tracking malicious activity for data protection reasons [41], let alone share any data with others. While it may be true that privacy regulations can limit the degree of detail in the data being shared, it should not be viewed as an impediment to sharing more aggregated data. Instead, in our view, claimed privacy restrictions are typically used as a cover for other disincentives to share data.

There can also be positive economic incentives for sharing security information. Gal-Or and Ghose developed a model of where sharing can work [70]: they argue that information sharing can encourage additional security investment. It is certainly true that the providers of security services stand to gain by sharing information, which can drive up demand.⁹ More generally, where there is a lack of industry awareness to threats, sharing information can certainly foster broader investment. This tendency to simultaneously share information and spend more on security has a more profound effect on highly competitive industries where product substitutability is higher. Where there is less competition it will be harder to get the market leader to share information. Gal-Or and Ghose also found that formal sharing organisations are more effective (in terms of

⁹In 2005 RSA Data Security bought Cyota, which runs a security data sharing scheme among US banks.

information sharing and investment spurred) when members join sequentially. By joining first, market-leading firms bootstrap the alliance by demonstrating their commitment to share information and invest in security, which encourages others to subsequently join.

While governments can specify requirements for data collection, it is up to the stakeholders to actually provide the data. Security vendors will feel it in their interest to provide inflated statistics; this has occurred frequently in the case of phishing. For example, the anti-phishing group PhishTank has boasted about the large number of sites it identifies [109], when in reality the number of duplicates reduces the overall number several fold. APACS, the UK payment association, provides another example by asserting a 726 % increase in phishing attacks between 2005 and 2006 (with merely a 44 % rise in losses) [8].

ISPs, by contrast, have an incentive to undercount the amount of wickedness emanating from their customers, particularly if they are held to account for it. But there is an even more pernicious problem with ISP reporting: ISPs hold important private information about the configuration of their own network that influences measurements. In particular, policies regarding dynamic IP address assignment can greatly skew an outside party's estimate of the number of compromised machines located at an ISP. ISPs also regard the size of their customer base as a company secret, which makes cross-ISP performance comparisons difficult.

Governments which endeavour to develop better information security statistics, that are based at least in part on private-sector data, must be aware of these biases and apply appropriate corrections and countermeasures. In an ideal world, mechanisms would be designed as strategy-proof, so that participants have no incentive to lie; in the real world, the statistician must strive to understand the application domain, set clear standards for data collection, and devise consistency checks across different sources and collection methods.

In more mature sectors of the economy, we can see useful examples of statistical institutions collecting business data jointly with industry bodies. For example, car registrations are handled by national authorities, before being aggregated to the European level by the Association of European Automobile Manufacturers,¹⁰ which then publishes Europe-wide figures. Safety and accident statistics for cars are collected by police and insurers. Another example comes from media circulation figures, which set the value of advertising space: in many Member States these figures are collected by private firms, some of them jointly owned and controlled by publishers and advertisers. Television ratings are collected by a panel of specially-equipped households. Best practice derived from these industries may inspire the evolution of accountable institutions that collect data relevant to information security.

At the behest of the European Commission, ENISA recently investigated whether to establish a framework for sharing collected data on information security indicators between interested parties [21]. They identified around 100 potential data sources, then surveyed a core of potential partners (CERTs, MSSPs, security vendors, etc.) who were invited to a workshop to further gauge interest. Unfortunately, there was very little desire for sharing raw data, aggregated data, or any information that doesn't already appear in the publicly-issued reports. The only initiative that received broad support is a 'high-

¹⁰<http://www.acea.be/>

level partnership’ called PISCE¹¹, which amounts to a wiki administered by ENISA linking to all of the reports, a closed mailing list, and possibly meeting occasionally to discuss their resources. Hence, mandatory reporting of particular indicators may be required for sharing to happen. Let us look now at the options.

4.3.2 Examples of information sharing

Option 1: Government-led ISACs across all of the CNI A US innovation was the exchange of data in closed industry groups known as information sharing and analysis centres (ISACs). Civil servants within the US federal government had worried about the protection of critical infrastructures (telecommunications, transport, water, chemical plants, banks, etc.) as these are mostly owned by private industry. Private firms have an incentive to under-invest in protection in the presence of externalities, and officials also worried about the growing dependence on the Internet. ISACs were set up as government-facilitated ‘talking shops’ in each critical industry for firms to share security-related information.

Their reception has been mixed. Early efforts centred on encouraging companies to establish ISACs within each sector. Some responded quickly, while others took several years to comply. Many firms were concerned about sharing security information with competitors and with the government [32]. We hear that many ISACs are moribund, and that other bodies have taken over de facto the information exchange role in many sectors.

Given this experience, it is hard to recommend that the EU follow the ISAC route. Instead we find it more prudent to examine particular types of data to be shared, then determine whether there are negative incentives that can be overcome through government assistance.

Option 2: Industry-led sharing Banks and other organisations targeted by phishing attacks have formed the Anti-Phishing Working Group (APWG) to fight the problem [7]; they have also created the more law-enforcement oriented (and more private) ‘Digital PhishNet’ organisation. The APWG shares information via regular closed meetings and by distributing a common feed of phishing URLs. Although it is based in the US, several European companies and banks participate. The push to create the APWG and to share information has been completely driven by the private sector. Most ‘take-down’ companies that provide outsourced phishing countermeasures are members. This is a very competitive sector, so we should not be surprised by the industry-led co-operation given Gal-Or and Ghose’s predictions.

Option 3: High-level partnership between data collectors The participants in a workshop organised by ENISA agreed to launch PISCE, a low-commitment, high-level partnership between organisations that publish reports on information security. ENISA serves as trusted mediator. Its initial goals are to ‘increase the visibility of existing data collections and mediate supply and demand’ using a wiki¹², categorise reports and facilitate understanding of reports without revealing details.

¹¹Partnership for ICT Security Incident and Consumer Confidence Information Exchange

¹²<http://wiki.enisa.europa.eu>

Perhaps PISCE will evolve into an arrangement where useful data is given to ENISA as input, who creates unbiased metrics. However, this is explicitly not part of its current mandate. Where voluntary data sharing is feasible, PISCE could perform a useful service. Where useful data is missing (e.g., from ISPs and the financial industry), mandatory sharing of specific data is the necessary outcome.

Option 4: Aggregated fraud figures driven by ENISA Elsewhere in the financial industry, the incentives against sharing security-related information must be overcome with additional regulatory encouragement. It is difficult for researchers and policymakers to prioritise spending on countermeasures when very little hard data about losses is available. Banks are often very keen to keep such figures private, for fear of repercussions in the stock price or scaring consumers.

One notable exception to this rule is the behaviour of APACS, the UK payments association, which has published aggregated figures for the annual amount lost to phishing attacks [8]. While the incentives are against individual financial institutions revealing losses publicly, a country-wide aggregation may still be useful to policymakers without inhibiting honest reporting very much.

We recommend below that ENISA should encourage similar financial-industry collections on a national and European level for different classes of online threats. Comparative national figures can be very helpful to a wide range of decision-makers given the differences in legal and technical approaches to fraud. For example, French banks have required PIN authentication for some time. When the UK was mulling over a switch to a similar technology, it would have been very useful if reliable, unbiased figures were easy to obtain.

Option 5: Network attack data-sharing with researchers One final area where information-sharing is important is at the IT level. Increasingly, Internet attacks require a global perspective for efficient detection and to understand attacker behaviour better. But companies naturally focus on the bit of the Internet visible to themselves.

Thus it would be ideal if ISPs could share relevant network-level information, whether directly or via third-party researchers. But there are significant impediments. First, ISPs are very hesitant to share any data that may reveal its network structure or the size of its customer base. Even if this information can be protected, sharing data on network traffic creates many privacy and legal complications. Much of the work of researchers investigating the econometrics of Internet crime is consumed in getting permissions for one set of data from (say) an ISP to be compared with another from (say) a mail service provider. While it would be helpful to make such research easier, it is not clear that systematic cross-ISP information sharing will be viable in the near future. However, statistics of the comparative performance of different ISPs are practical to collect, and they could provide a useful and powerful market signal.

4.4 Information sharing recommendations

Our recommendation is that ENISA's information sharing efforts should focus on industries with a clear benefit but where sharing is not already taking place in every Member State – and the two industries where more information should be made available are the financial industry and ISPs.

As noted above, the UK banks do present annual aggregate figures for fraud, via the Association of Payment and Clearing Services (APACS). As far as we have been able to determine, no other Member State publishes statistics of this kind. As banks collect such statistics for operational, internal control and audit purposes, and aggregating them nationally is straightforward, we believe this practice should become standard practice in the EU. The statistics are particularly critical to the formulation of policy on network and information security since the majority of the actual harm that accrues is financial. Without a good measure of this, other figures – whether of vulnerabilities, patches, botnets, or bad traffic – lack a properly grounded connection to the real economy.

Recommendation 2: We recommend that the Commission (or the European Central Bank) regulate to ensure the publication of robust loss statistics for electronic crime.

In many cases, fraud statistics are already collected by the police or banking associations, so regulatory action should aim at harmonisation of definitions, metrics and release cycles across Member States. A good first step would be to require figures broken down broadly as the APACS statistics are at present and show losses due to debit and credit card fraud (subdivided into the useful categories such as card cloning versus cardholder-not-present, national versus international, and so on).

It must be said that the UK regime is not perfect. For example, the UK government ordered that from April 2007 bank fraud should no longer be reported to the police, but in the first instance to banks – who would save police time by consolidating the reports and passing on details of those cases that they wanted prosecuted and for which they saw some chance of success. The effect is that reports of online fraud and card fraud have dropped to zero for many police forces. In addition, the UK system hides the identity of individual banks; although it's known that one particular bank suffered most of the phishing losses in 2006, the identity of that bank was not published by APACS. A senior bank official even remarked to one of us that they don't keep detailed records of complaints by social indicators, and thus have no way of telling if dispute resolution mechanisms discriminate against the less educated, or against customers from other Member States. Banks might resist being required to collect data that they don't already collect internally, but legislators might feel that issues of discrimination, access and the Single Market justify overruling the banks on this. Finally, some data relevant to the analysis of bank fraud may be collected from nonbank sources, such as telcos and anti-virus companies who may be in a better position to monitor the frequency of specific fraud vectors (such as phishing versus keyloggers).

As for the information that should be published by and about ISPs, it is well known at present within the industry that some ISPs are very much better than others at detecting abuse and responding to complaints of abuse by others. This is particularly noticeable in the case of spam. A small-to-medium sized ISPs may find its peering arrangements under threat if it becomes a high-volume source of spam, so such ISPs have an incentive to detect when their customers' machines are infected and recruited into botnets. A typical detection mechanism is to look for machines that are sending email directly, rather than via the ISP's smarthost facility; infected machines can then be placed on a subnet that gives them restricted access to the Internet, so that they are able to access anti-virus software and have low-bandwidth connection to random websites, but where a firewall

stops them sending spam while their owners are encouraged to clean them up. Large ISPs don't face the same peering-arrangement pressures, so as a result some send significantly larger quantities of spam and other bad traffic than others. We feel it would be strongly in the public interest for quantitative data on ISPs' security performance to be available to the public.

Recommendation 3: We recommend that ENISA collect and publish data about the quantity of spam and other bad traffic emitted by European ISPs.

As Europe has some 40,000 ISPs, a staged approach may be advisable – with initial reports collected using sampling, followed if need be by action through telecomms regulators to collect more detailed statistics. However, even rough sample data will be useful, as it's the actions of the largest ISPs that have the greatest effect on the level of pollution in the digital environment.

Anyway, we feel that ENISA should take the lead in establishing these security metrics by setting clear guidelines, collating data from ISPs and other third parties, and disseminating the reported information. To begin with, ENISA could make a positive contribution by collecting and disseminating data on the rate at which ISPs are emitting bad packets. Such data could serve as a useful input to existing interconnection markets between ISPs since high levels of bad traffic can be costly for a receiving ISP to deal with.

The types of digital pollution to be measured must be defined carefully. To track spam, useful metrics might include: the number of spam messages sent from an ISP's customers; the number of outgoing spam messages blocked by an ISP; the number and source of incoming spam messages received by an ISP; and the number of customer machines observed to be transmitting spam for a particular duration. To track other types of malware, the number of infected customer machines would be relevant, along with the duration of infection.

Once data are available on which ISPs are the largest polluters, the next question is what should be done about them. This moves us from the heading of 'information asymmetries' to our next heading, 'externalities'.

5 Externalities

As noted above, externalities are the side-effects that economic transactions have on third parties. Just as a factory belching out smoke into the environment creates a negative externality for people downwind – and indeed for the whole world in the case of global warming – so also people who connect infected PCs to the Internet create negative externalities in that their machines may emit spam, host phishing sites and distribute illegal content such as crimeware.

5.1 Fixing externalities using carrots

Subsidy is one of the traditional (supply-side) policy instruments for dealing with externalities. The EU's Framework Programmes of research have not only been used to develop many technologies that deal with environmental pollution, but also to develop many security technologies. The most notable is probably the smartcard industry.

Europe dominates the smartcard business; according to a recent market survey, card sales are currently USD 2.3 billion worldwide. Prices vary from USD 0.4 to USD 2; the mean price may be about USD 1. Smartcards are widely used in mobile phones as SIM cards, in pay-TV as subscriber cards, and in banking with the EMV protocols. The value chain now includes not only hardware designers such as ARM, foundries such as Infineon and OEMs such as Gemplus, but also a wide industry of terminal makers, testing labs, and specialist software vendors producing everything from SIM toolkits to back-end software for bank card systems. In contrast, in the US where the smartcard industry has not received comparable government support, bank cards still predominantly use old-fashioned magnetic strips (although some banks are starting to introduce RFID-based smartcards). The smartcard industry is thus viewed as the poster case of economic development in the technology field being spurred by state intervention. There are many less well-known examples, such as the 'Bolero' system for cryptographically-secured electronic bills of lading that facilitates trade within the EU and globally.

Buying innovation The other traditional (demand-side) policy instrument is to use public-sector purchasing. Outside the European Union, the development of multilevel secure (MLS) systems was driven by the US Department of Defence for over a quarter of a century. MLS products enforce access-control rules relating to information classification – for example, that a user cleared to SECRET is not allowed to see information classified at TOP SECRET – and this enforcement is independent of user actions. The main fruits of this purchasing program are, first, Trusted Solaris, the high-security version of Sun's operating system that is now included in the standard Solaris distribution and is thus also available to private buyers; second, the SELinux version of the Linux operating system, developed with assistance from the NSA, that is freely available and is incorporated in the Red Hat distribution; and third, the mandatory access control features now being shipped by Microsoft in their Vista operating system. This is a good example of how public procurement stimulates the development of a product that might not have been competitive if private purchasers had had to bear the fixed development costs.

Purchasing innovations is different from purchasing commodities in a specific user-producer interaction during the contract negotiation phase: potential suppliers must learn

about the procurer's needs and the suppliers' knowledge of possible technical solutions must be passed back to the procurer [42]. Conducting this interaction in a fair and transparent manner during a public tender with various suppliers is acknowledged to be challenging, but doable. In particular, it is indispensable that the procuring authority has very good technical knowledge. A 2005 report to the European Commission concludes from a systematic country overview: 'The only EU Member State¹³, which has started a broad strategic process for the usage of public procurement to foster innovation, is the United Kingdom' [68].

Buying assurance Another way in which public-sector bodies can use their purchasing power to enhance information security is by purchasing assurance. The prominent example of this at present is the Common Criteria, a scheme for evaluating information security products which is jointly run by thirteen Member States, along with the US, Canada, Australia, New Zealand, Japan, Singapore, Turkey, India, Israel, Korea and Malaysia. The Common Criteria provide a framework within which firms can have their products tested at approved laboratories, and have evaluations recognised across participating countries for the purposes of government procurement. Other industries have bought into the Criteria; for example, VISA is starting to use Common Criteria evaluations rather than its own evaluations for the PIN entry devices used in the EMV payment protocol for bank smartcards ('chip and PIN'). A further, but less prominent, example is the establishment in the Netherlands and Sweden of laboratories that evaluate clinical information systems – not just for security, but also for safety and interoperability – and provide an 'approved products list' for doctors and hospitals in their countries. Such a scheme was recommended for the UK also by a recent parliamentary inquiry there. And as we shift the focus from security to the broader question of safety, there are numerous standards and arrangements in specific industries (burglar alarms, cars, aircraft, electrical goods, ...).

So there is plenty of precedent for the public sector to use its purchasing power, whether directly or indirectly, to improve the state of system security (and safety). As well as evaluation being performed by laboratories licensed by the government (as with the Common Criteria), it can be done by insurance laboratories (as with burglar alarms) or on a basis of self-certification by vendors. Self-certification should bring with it some penalty mechanism: for example, if a manufacturer warrants that his product is safe (or secure) and it turns out not to be, then he should be liable for damage (we will return to this when we discuss liability).

Although no reliable data is available on the consolidated volume of EU-wide public purchases of IT, the aggregated buying power of administrative bodies, healthcare systems and educational institutions should be too large to be ignored by manufacturers. The European procurement framework as established with Directives 2004/17/EC and 2004/18/EC leaves room for innovation oriented procurement [68]. There are thus significant opportunities to use procurement as a strategic tool to lift barriers to network information security.

¹³EU15

5.2 Fixing externalities using sticks

As far as security externalities go, the volume issue is malware that's used to harm others, rather than the infected host. At present, such malware is the backbone of the underground economy in electronic crime. It can be used to send spam, host illicit sites for phishing and hawking shady goods, launch denial of service attacks, and even search for more vulnerable hosts to infect.

Such malware is installed using social engineering; using weaknesses in core platforms – operating systems, communications systems (e.g., routers) and server software; or increasingly by exploiting applications. The incentives are not as misaligned for core platforms – Microsoft has been improving its security for some time, for example, and stands to suffer in terms of negative publicity when undisclosed vulnerabilities are publicised.

However, exploits at the application level will need a different approach. Users readily install add-on features to web browsers, enable web applications run by untrustworthy firms, and run unpatched or out-of-date software. They may also choose not to install or update anti-virus software.

5.2.1 Control points

There are a number of *control points* where we might possibly do something about system insecurity. We discussed the exploit lifecycle in Section 2.3: vendors carelessly introduce vulnerabilities; people discover them; vendors fix them; they nonetheless get exploited for a while; machines get recruited to botnets; they are discovered to be infected; they are removed from the network for disinfection; and the stolen assets are recovered.

We have discussed the incentives facing vendors, and what can be done about them using the carrot of public purchasing. In Section 6 we will discuss what can be done with the stick of liability: there have been repeated calls for software and platform vendors, as well as service providers, to be held responsible for the damage caused by the bugs in their systems. This is likely to be part of the solution, but it is unlikely to be the whole solution since attacks are often due to poor configuration and late patching.

The next influential control point is the ISP. ISPs control a machine's Internet connection, and therefore its ability to harm others. There are many steps an ISP can take to limit the impact of malware-infected customer devices onto others, from disconnection to traffic filtering.

The machine owner is another important control point. Large companies manage their machines in several ways. First, they have a network perimeter where they can deploy devices such as firewalls to minimise exposure to compromise as well as restrict outbound communications from compromised machines. They also employ technicians to repair infected devices.

For regular end users and SMEs, there are fewer steps that can be taken. One is to maintain updated software, from the OS to applications and anti-virus tools. However, users cannot protect themselves at the network perimeter as effectively as large businesses can, and furthermore they can have tremendous difficulty repairing compromised devices.

A useful analogy, to which we'll return later, is road safety. The incidence of injury-causing road traffic accidents in developed countries is now less than a tenth of the rate in less-developed countries such as China, or indeed in Europe and America between the wars. The improvement is due to a number of factors: cars are safer (as manufacturers are

now liable for defects, and crash test ratings of cars are published); roads are much safer, with uniform standards for construction, lighting, signage and crash barriers; drivers are better trained; tachographs restrict commercial drivers' working hours; police forces arrest drunk drivers; and cultural change has made drunk driving socially unacceptable. This is not because of the damage that the drunk does to himself, but because of the externality – the harm the drunk does to others.

Infected machines are the main source of harm to others, and many of them are not running current antivirus or properly patched software. However, at this point the analogy with road traffic becomes somewhat strained. Many infected machines do have antivirus software, as the more competent malware writers test their products carefully against existing antivirus products, and often users fail to patch for apparently good reasons (see Section 6.5). In some cases, the consequences of such failures should really be the liability of the vendor rather than the end-user – a point to which we will return later.

Anyway, machines get infected. Where the machine belongs to a large company, there are professional staff to detect this and so the clean-up; but for lone users and SMEs, the task falls to either the user or the ISP.

Compared to the other stakeholders, ISPs are in the best position to improve the security of end-user and SME machines. They control user access to the Internet; they can implement egress filtering to limit the impact of compromised machines on others; they are well-positioned to carry out network-level tests of system security; and they have the ability to communicate with their users by telephone or postal mail, not just by Internet channels. As relatively large organisations, ISPs can also realise economies of scale not possible for SMEs and end-users.

ISPs are divided on whether they should actively isolate infected customer machines, let alone whether they should take active steps to prevent infection. An Arbor Networks survey found that 41% of ISP respondents believed that they should clean up infected hosts, with 30% disagreeing and 29% uncertain [96]. Taking costly steps to repair customer machines, potentially including the unpopular move of temporarily cutting off service, is undesirable for ISPs when most of the negative effects are not borne by the ISP. Yet, as noted, a number of well-run ISPs do take suitable measures, such as confining infected machines to a filtered subnet, because of the direct and indirect costs to an ISP of becoming a source of digital pollution.

5.2.2 Policy options for coping with externalities

So if ISPs should take actions to raise the level of end-user security, then what is the best policy option to encourage them? We discuss and evaluate several options: exhortation via best practices, taxation of observed bad emissions, a cap-and-trade system, liability assignment, and fixed penalties.

Prerequisite: Publish data on ISP performance A key prerequisite for every policy option just discussed is identifying consistent metrics of malware. It is well known in the industry that some ISPs have many more infected machines than others, and send vastly greater amounts of spam and service-denial traffic, but there's a shortage of public numbers. The prevalence of 'dynamic' allocation of IP addresses makes it hard to understand the statistics and means that third-party researchers cannot fill the gap. We already re-

commend that ENISA collect and publish data about the quantity of spam and other bad traffic emitted by European ISPs (Recommendation 3). A public league table can of itself raise the overall performance of ISPs by highlighting both over- and under-achievers, as has been the case for league tables published in other areas. For example, the UK government now publishes an annual car theft index, which identifies the makes and models of cars alongside the number of reported thefts. While causality is difficult to prove, the overall number of car thefts in the UK has dropped considerably since the index's introduction.

Once decent statistics are available, the next question is what further incentives might help. Given car-theft statistics, for example, governments could have left it to the car industry to make their cars harder to steal; this was the policy pursued for many years, but with little effect (the car makers kept on producing vehicles using simple mechanical locks that could be bypassed easily). In the end it took the fall of the iron curtain, which led to a surge in car crime, which in turn led German insurers to pressure car makers to fit new vehicles with remote key entry devices using cryptographic authentication with the engine control unit. This dramatically cut car theft throughout Europe, and probably contributed to falls in vehicle-borne property crime as well.

How might online crime be tackled? The first option is *laissez-faire*.

Option 1: Encouraging self-regulation, perhaps with the threat of intervention The most hands-off response is to use indirect regulatory tools such as encouraging ISPs to adopt best practices through self-regulation. Once coherent statistics comparing ISP performance are available, governments could pressure under-performing ISPs into improving their behavior.

However, there are several reasons to doubt the efficacy of self-regulation. The foremost reason is that it has not worked very well so far. An OECD report investigating the economics of malware [41] argues that some positive incentives exist for ISPs to take precautionary measures. It points to the high cost of customer support, since consumers often call their ISP when they experience problems with their computer. However, this incentive does not work when the malware is designed to be undetectable, as is the case for botnets. Other incentives put forward include the threat of address blacklisting and reputation benefits.

Indeed, some ISPs have taken action, most likely in response to these (weak) positive incentives. However, the poor performance of other ISPs overshadows this. This is because overall Internet security is down to the 'weakest link': attackers identify and exploit the worst-performing ISPs. If one ISP takes steps to clean up infected machines, the attackers may simply compromise more machines at a less diligent ISP. Worse, the incentives for improvement fall largely on the smaller ISPs, whose peering arrangements may be at risk if they send too much spam, and not on the larger ones. Hence vast quantities of digital pollution emanate from a number of large ISPs, who face limited economic incentive to clean up their act.

Option 2: Taxing 'digital pollution' Taxation is a traditional policy tool. ISPs that fail to take suitable measures to prevent their customers becoming a nuisance to other Internet users could be taxed directly. Taxation is likely to be vehemently resisted by ISPs, who will say that it is their customers causing the problems (a stakeholder at

our consultative meeting inveighed against ‘punishing the innocent’). However, taxing customers directly would be even more controversial than indirectly doing so by taxing the ISPs. One further difficulty with taxation is setting an optimum rate of tax.

Option 3: A cap-and-trade system An alternative to direct taxation is a cap-and-trade system, as used already with carbon credits. Under such a scheme, ISPs would either have to install proper filtering systems, or purchase ‘emission credits’ from other ISPs that had done so. In theory, trading schemes enable firms to reduce their emissions (whether of carbon or of wickedness) at the lowest possible cost. The carbon experience has shown some practical problems – with the allocation of initial rights, the definition of reliable metrics for the amount of ‘pollution’, and possible regulatory arbitrage.

There are reasons to be optimistic about the prospects of a cap-and-trade system for ‘digital pollution’. Because there is great variation in the size of ISPs, many smaller providers might prefer to avoid the capital costs of good filtering. For them, it may be cheaper to buy credits off larger providers that have implemented industrial-scale filtering anyway for other reasons (for example, to block child pornography). Such a market could also make filtering technologies more attractive to mid-level ISPs who do not find filtering cost-effective at present.

However, a cap-and-trade system must still overcome the other pitfalls experienced by the carbon-trading system. First, it is suboptimal to provide extensive and permanent rights to pollute for free. Consistent metrics are especially important, since organisations will be trading on the measurements. At present, only spam can be measured in a universally-recognised manner. Other potential pollution types, such as malware incidents, phishing sites or denial-of-service attacks, cannot be measured in a consistent way across the industry at present. Another problem is the unpredictability of pollution levels. Power companies know how much carbon is emitted from a coal-fired plant and purchase credits in advance. For ISPs, the situation would be more complicated because the pollution levels depend on whether they are targeted by attackers. These issues argue against a cap-and-trade system, at least for the present.

Option 4: Assigning liability of infected customers to ISPs Externalities might be dealt with through liability assignment. Legislation could allow any party that suffered harm to sue an ISP whose customers had connected malicious machines to the Internet. It is essential that liability be placed on ISPs and not consumers, since ISPs are in a position to take remedial steps. The ISPs will doubtless claim that they are unaware of the malicious machine and that they are unable to prevent the harm. Their failure to respond to notification or to invest in suitable blocking equipment is something they can easily fix. However, there are two rather more serious difficulties with imposing liability in quite this manner. First is the potentially high transaction cost of lawsuits. Second is the difficulty of valuing the monetary loss associated with individual events.

Option 5: Fixed-penalty charges for ISP inaction To deal with both the uncertain costs of liability and the difficulty for users of proving a quantum of damages, another option is to instead introduce fixed penalty charges if ISPs do not take remedial action within a short time period of notification. There is great variation in the response times

for ISPs when notified that their customer's machine is infected. At present, the best-performing ISPs can remove phishing sites in less than one hour, but some ISPs take many days or even weeks to respond. Introducing a fixed penalty for machines that continue to misbehave after a reasonable duration, say 3 hours, would drastically speed up remedial action.

Fixed penalties are useful because they avoid the problem of quantifying losses following every infringement. They have been used effectively in the airline industry, where the EU has introduced penalties for airlines that deny passengers boarding due to overbooking, cancellations or excessive delays. The goal of this regulation is provide an effective deterrent to the airlines. Fixed penalties are also routinely used for traffic violations. Again, the penalties deter violations while simplifying the liability when violations occur. The threat of penalties should alter behavior so that, in practice, fixed penalties are rarely issued.

For fixed penalties to work, a consistent reporting mechanism is important. Fortunately, existing channels can be leveraged. At present, several specialist security companies already track bad machines and notify ISPs to request their removal. This process could be formalised into a removal notice. End users should also be allowed to send notifications. For example, if a user receives a spam email, he could send a notification to `abuse@isp.com`, as is already possible.

One issue to consider is to whom the fixed penalty should be paid. To encourage reporting, the penalty should be paid to whoever sent the notice. What about duplicate payments? One compromised machine might, for example, send millions of spam emails. If a fixed penalty had to be paid for each received report, then the fine may grow unreasonably large. Instead, the penalty should be paid to the first person to report an infected machine, or perhaps to the first ten who file reports. (Enabling legislation should leave enough room for the scheme to be modified in the light of experience.)

Issues of proportionality and possible side effects: Given the threat of stiff penalties for slow responses, ISPs might become overzealous in removing reported sites without first confirming the accuracy of reports. This might lead to a denial-of-service-attack where a malicious user falsely accuses other customers of misdeeds. There is also the established problem that firms who want machines taken down for other reasons – because they claim that it hosts copyright-infringing material, or material defamatory of their products – are often very aggressive and indiscriminate about issuing take-down notices. These notices may be generated by poorly-written automatic scripts, and result in risk-averse ISPs taking down innocuous content.

In theory, a user can tell her ISP to put back disputed content and assume liability for it, but often the ISP will then simply terminate her service, rather than risk getting embroiled in a legal dispute. This is likely to become a serious problem, and it's not helped by support from the President of France for a music-industry initiative to disconnect file-sharers from the Internet. We believe that this initiative will founder, because network connectivity is as important nowadays as a supply of water or electricity; without it a house is for many people uninhabitable. The problem remains that, in many countries, ISPs have got into the habit of writing their contracts so that they can terminate service on no notice and for no reason; a regulatory intervention will probably be required here, and we will discuss it later.

For now, let us discuss the less politically-charged problem of how to remove infected machines, as opposed to machines whose content is the subject of legal dispute. There are two options. First, reporters should be held liable for the accuracy of their accusations – so anonymous submissions should be disallowed. Second, there has to be a ‘put-back’ mechanism that users can invoke to get their ISPs to reconnect an incorrectly classified machine quickly. Even given these measures, a penalty system might be abused by competitors (as with click fraud for pay-per-view ads), and in any case many business users will be unwilling to take the risk of being disconnected by their ISP following an allegation of infection.

Another necessary precaution is ensuring that the ISP does not automatically shift the penalty on to the consumer whose machine triggered the penalty. We suggest that regulation provide for two classes of contract – consumer contracts that limit consumer liability (say to EUR 50), and business contracts that can be written as the parties see fit. Many businesses will prefer to take over responsibility for machine clean-up, and liability for penalty charges, from their ISP in return for assured continuity of service. Small businesses may not have the capability to clean up machines, but will still want assured service: they can contract out their security management to third parties. (This is a growing sector, and one in which some large ISPs are already active; for example, BT bought Counterpane in 2006.)

It is not the purpose of this report to undertake the detailed design of a fixed-penalty system, as this would have to evolve over time in any case. We nonetheless feel that it is the single measure most likely to be effective in motivating the less well-managed ISPs to adopt the practices of the best.

Recommendation 4: We recommend that the European Union introduce a statutory scale of damages against ISPs that do not respond promptly to requests for the removal of compromised machines, coupled with a right for users to have disconnected machines reconnected by assuming full liability.

We understand from the stakeholders’ meeting that this is the most controversial of our recommendations. We therefore say to the ISP industry: do you accept it’s a problem that infected machines remain connected to the Internet, participating in botnets for extended periods of time? And if so, what alternative means do you propose for dealing with it?

To return to our road-safety analogy, the operator of a highway cannot reasonably claim that all the responsibility for road safety must fall on the car makers and drivers; the combination of modern cars but poor roads has been lethal in China. Stretching the analogy, if one particular intersection pours large numbers of drunk drivers on to the highway every Friday night, who’s to stop that? We patrol physical highways using police officers employed by the government; do we also need policemen in each ISP dealing with infected machines, or could the ISPs’ own staff do it more efficiently and cheaply?

6 Liability assignment

Liability raises much broader issues than just whether ISPs should be liable for not taking down infected machines promptly. One issue that has been raised repeatedly over the years is whether software vendors who sell insecure products should be liable for the harm that they cause. It is widely believed that the aggressive liability disclaimers found on almost all software license agreements protect vendors from lawsuits.

6.1 Analogy with car safety

There is an interesting analogy between online safety and automobile safety. For the first sixty years of its existence, the car industry managed to avoid most of the liability for design and manufacturing defects. Vehicles were not equipped with seat belts or crumple zones, as the vendors considered aesthetics more of a selling point than safety. Eventually public opinion changed, catalysed by Ralph Nader's book *Unsafe at Any Speed* [104], and by US case law enabling accident victims to sue the manufacturer and not just the driver or the car dealer [115]. This led to a change of attitude by car makers, helped along by a multitude of regulatory interventions, ranging from the publication by government of 'star ratings' for vehicle crashworthiness to specific laws and regulations on seatbelts, airbags, etc. Vehicle-safety initiatives were complemented by driver training, the construction of highway systems, and the steady improvement of road standards for lighting, signage, and crash barriers. The rate of injury-causing accidents in the US and Europe is now more than an order of magnitude less than in China where the implementation of this package of measures has been patchy at best.

Given that the first software was written in 1949, and the first software was sold sometime in the 1950s, we are now getting to a comparable point in the software industry's evolution. It is also becoming clear that as our civilisation comes to depend more and more on software, the culture of impunity among software writers and vendors cannot continue indefinitely. For example, the UK House of Lords Science and Technology Committee recommended, in the context of an inquiry into Personal Internet Security, that that the UK government should, working through the EU, seek to rectify the inappropriate liability assignments in the medium term [76]. Similar comments have been made in respect of the ability of key service providers (from telecomms to electronic banking) to disclaim liability for failures and frauds. A special adviser to President Bush remarked in 2004 that it was unsustainable to hold software companies blameless, and hoped that liability would be fixed by the courts as the only institution with the flexibility to adapt to rapid technological change [115]. It is a long-established principle in tort law that one should assign liability to the party best able to prevent the undesired outcome. Security economics research underlines this: if a system vendor or operator can dump liability, he will make suboptimal effort.

Many people have argued specifically that if Microsoft were liable for the consequences of the many exploits of Windows, then the company would invest much more heavily in securing it. Microsoft for their part will argue that in recent years they have made enormous investments in making Windows more secure (and there is much truth in this). It is also argued that a move to software liability could be harmful to the free software community: graduate students would be much less willing to contribute code to the Linux

project if they faced the prospect of being sued in later years if a bug they introduced allowed a critical system to be compromised. A reluctance to embrace software liability is thus one of the few issues that unites the proprietary and free-software worlds. Microsoft further argued, at our consultative meeting, that it would be unfair to impose liability for software but not services: for example, the functionality provided by one of their top-selling software products (Office) is also provided by Google as an advertising-supported online service (Google Documents). This is a valid point, and we will return to it below.

Software (and service) liability is a huge and complex issue, just like automobile safety was in the 1960s. It is unlikely to be fixed by a single over-arching Directive that assigns liability unequivocally to vendors, any more than car safety was. An attempt to make Microsoft liable for all the harm caused by vulnerabilities in its systems would be strongly resisted, not just by the company but by the US Government. It would raise many broader issues. The history of the twentieth century teaches that it's a bad idea for governments to intervene in private contracts without good reason; markets are generally more efficient and the main justification for regulatory intervention is market failure. Examples of grounds for intervention are the protection of consumers, who frequently lack both information and the bargaining power, and monopoly. We will return to consumer protection shortly.

6.2 Competition policy

As for monopoly, had this report been written two years ago, we might have been concerned about the dominance of Cisco in the router market and Symbian in the market for mobile-phone operating systems. As Cisco sold most of the routers used in the Internet backbone, there was a potential critical-infrastructure vulnerability: if a flash worm had come round that damaged Cisco equipment, the Internet backbone could have been taken down, causing considerable economic damage. However, recently Cisco's prices have evoked competition from Juniper and others, so that the situation is improving.

In general, contracts work fine for businesses where there's competition, and so it would seem to be reasonable at this time to deal with liability issues on a sectoral basis. Europe has more competitive communications service providers than the US (hence network neutrality isn't as acute an issue here as it is there) but more concentrated financial services (the US Glass-Steagall Act of the 1930s left America with many small banks rather than the handful of large ones in a typical Member State). For example – moving from software liability to systems liability – one problem in several Member States is that banks don't compete very vigorously to acquire credit card transactions from merchants; in the UK, which has only three large acquirers, there have been repeated findings by the competition authorities against the banking industry in this regard (see for example [107]). The effects of this are both financial (merchants pay more to process credit-card transactions) and on liability (UK banks make merchants liable for cardholder-not-present transactions). This risk dumping may be partly to blame for the continuing rise in online fraud against UK cardholders; we hope that once comparable figures are available across Europe, policy effects like this will become clearer. However, despite its relevance for online crime, the contracts between banks and merchants are fundamentally a matter for financial and competition-policy authorities.

It was argued at the stakeholders' meeting that vendor liability would constrain inter-

operability and thus weaken competition. We do not believe this. The many monopolies and imperfect competition in the software market arise from the well-understood phenomena of network externalities and lock-in, and the lack of interoperability is usually quite deliberate [124]. From a technical point of view, there is no reason to believe that secure and interoperable designs are mutually exclusive.

6.3 Product liability

Contrary to popular belief, software vendors and service providers are not in a position to lawfully disclaim all liability to their customers, and of course their customer contracts have little effect on their liability towards third parties. Returning to the car safety analogy, a car maker might sign a contract with a customer saying that the maker would not be liable to the customer for injury, but if the steering fails and the car injures a third party who has not signed this contract, then that third party can sue. But who should they sue? For years, the car makers argued that they should sue the driver at fault, who in turn would sue the person from whom he bought the car if he believed that the cause of the accident was a design defect rather than his own negligence. That person in turn might sue the person from whom he bought the car, and so on, until eventually a lawsuit arrived at the car maker's factory. Needless to say, this placed an enormous burden on the victim of a design defect, and, following a series of court cases from 1916, the US courts eventually ruled in the landmark *Greenman v. Yuba Power Products* case in 1963 that the victim of a design or manufacturing defect could sue the maker of the defective product directly [115]. This principle arrived in European law in 1985 via the Product Liability Directive which adopted language similar to that used by Judge Traynor in *Greenman* [49].

This Directive makes software vendors liable, despite what their contracts may say, for personal injury and property damage caused by design defects. It's unclear whether software is always covered, though it often will be: the Directive states

Article 1

The producer shall be liable for damage caused by a defect in his product.

Article 2

For the purpose of this Directive, 'product' means all moveables even if incorporated into another moveable or into an immovable. 'Product' includes electricity.

It would thus appear that if a citizen buys a copy of Office in a shop, installs it on a PC, and suffers injury or property damage as a result (for example) of a bug in Word or Excel causing him to make an incorrect tax filing, he can already sue. Indeed, in the UK, such a case has been brought successfully under a different legal theory, unfair contracts, which we will describe below; and there are many laws at the national

level under which a software vendor could be sued by an injured customer regardless of contract disclaimers. (In the UK, for example, the vendor can be liable for common law negligence, or under the Misrepresentation Act 1967, or the Sale of Goods Act 1979, or the Sale of Goods Supply of Services Act 1982.) There remain some unclear points of law, for example whether a software sale is the sale of a good or the supply of a service; and the position may be affected by whether the copy of Office were bought in a physical box or as an electronic download. We'll return to such matters later. The point we make here is that the alleged 'immunity' of software vendors is in fact a myth. (It may be effective in dissuading people from litigation, but it is a myth nonetheless.)

However, as software becomes embedded in more and more devices on which we rely in our daily lives, Microsoft Office is not perhaps the best motivating example. A better one (for which we thank Alan Cox) is a navigation system. Suppose that a truck driver purchases a navigation system and, relying on it, is directed by a software error down a small country lane where his lorry gets stuck, as a result of which a valuable load of seafood is spoiled. This case is interesting because navigation can be supplied in a number of ways as a product, as a service, or as a combination of both.

1. A common way to get a navigation system is to buy a self-contained GPS unit in a shop.
2. A driver can also get a navigation system in the form of software to run on his PDA or laptop computer.
3. Navigation is also available as a service, for example from Google Maps.
4. An increasing number of high-end mobile phones have built-in GPS, and can also provide route advice either through embedded software or an online service.
5. The driver could hook up the GPS receiver in his mobile phone to route finding software in his laptop.
6. As well as proprietary route-finding systems, there's a project¹⁴ to build a public-domain map of the whole world from GPS traces submitted by volunteers. In addition, a driver's proprietary system might run on an open platform such as Linux.

So the question is, which of the above suppliers could the truck driver sue? Certainly it's common for GPS equipment vendors to put up disclaimers that the driver has to click away on power-up, but the Product Liability Directive should deal with those. This suggests that we should be able to deal with the liability issues relating to embedded systems – that is, the software inside cars, consumer electronics and other stand-alone devices – as a product-liability matter. (Consumer law – and in particular the Unfair Contract Terms Directive – reinforces this; we'll come to it shortly.)

¹⁴See <http://www.openstreetmap.org>

6.4 Software and systems liability options

Following this discussion, we conclude that software liability is a large and complex issue, and one that's widely misunderstood. Clearly something needs to be done about it; our civilisation is becoming ever more dependent on software, and yet the liability for failure is largely disclaimed and certainly misallocated. What are the options?

Option 1: Make the vendors liable The big-bang approach would be a Directive rendering void all contract terms whereby a software vendor or system supplier disclaims liability for defects. It is likely that this option, however fervently sought by the more outspoken critics of the software industry, would be bad policy. As discussed above, governments should not interfere in freedom to contract unless they have good reason to; and there is merit in the Microsoft point that software should not be singled out for unfair and discriminatory treatment. In addition, a 'Software Liability Directive' would probably not be politically feasible because of the vigorous resistance it would provoke from all across the software industry and indeed from the US government.

We believe that, as with the motor industry, a patient and staged approach will be necessary. While it might have been feasible to impose stricter rules on software liability as late as the 1970s or even 1980s, by now there is software in too many products and services for a one-size-fits-all approach to be practical. In particular, where software from dozens of vendors is integrated into a single consumer product, such as a car, the sensible approach (taken by current EU law) is to hold the car maker (or primary importer) liable for faults that cause harm; this ensures that the maker has the right incentives to ensure that the software in their product is fit for purpose. Thus, for the time being at least, liability for failures of software in embedded systems should continue to rest with the maker or importer and be dealt with by safety, product-liability and consumer regulation.

However, where devices are connected to a network, they can cause harm to others. Cyber-criminals can in principle use any network-attached device – be it a PC, a mobile phone, or even a medical device – to launch service-denial attacks, send spam, and host unlawful content such as phishing websites and indecent images of children. A case has been made, for example, that US lawmakers should create a specific tort of the negligent enablement of cybercrime [115]. Even if the EU is not going to have a 'Software Liability Directive', does it need a regulation creating liability for vendors who negligently put into circulation large numbers of devices that are easily infected by crimeware?

Option 2: More specific rights to sue for damages If our fourth recommendation, namely that there should be fixed-penalty charges on ISPs who fail to take down infected machines promptly once put on notice, is accepted, then there would be a case for the ISPs to be able to recover some or all of these charges from the responsible parties. As we noted above, it is advisable to limit the amounts that can be recovered from individual consumers, and so it is logical to enable ISPs to recover their charges and costs from software vendors who negligently supply vulnerable software to consumers. It may indeed already be the case that ISPs could already sue Microsoft and prevail using national transpositions of the Product Liability Directive. One might argue in favour of regulatory action to make this point clear. However this is a somewhat indirect way of proceeding; we might have to wait years for an ISP or other injured party to drum up the courage to

launch the needed test case.

Option 3: Laissez-faire The third option, which should at least be mentioned, is to do nothing. For example – as we will discuss below – Sun and Hewlett-Packard are much slower to patch than Microsoft or Red Hat, and so (in the business sector at least) the mere provision of authoritative, unbiased information about the level of assurance provided by different vendors’ offerings may be sufficient to enable competitive pressures to fix the problems over the medium term. In the case of consumers, however, there is little choice: people can either buy Windows, or pay significantly more for Apple machines (which also run fewer applications).

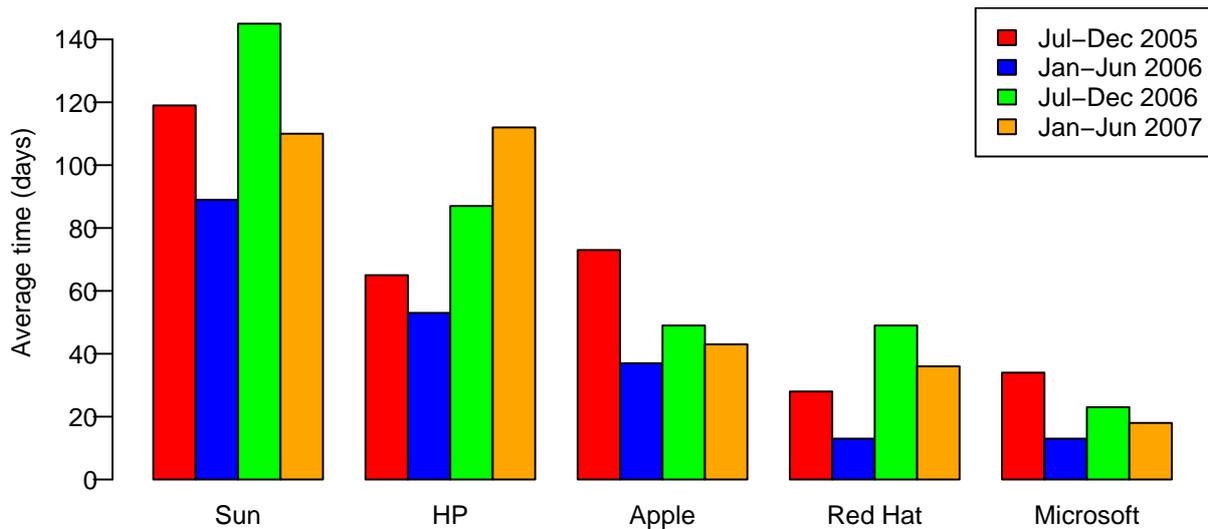
Option 4: Safety by default The fourth option is that, when selling PCs and other network-connected programmable devices to consumers, vendors should be required to configure them so that they are secure by default. It’s illegal to sell a car without a seatbelt, so why should shops be allowed to sell a PC that doesn’t have an up-to-date operating system and a patching service switched on by default? We believe that this gives a more direct approach to the problem than option 2; and of course vendors who sell insecure systems should be exposed to lawsuits from ISPs and other affected parties.

Recommendation 5: We recommend that the EU develop and enforce standards for network-connected equipment to be secure by default.

The precise nature of ‘secure by default’ will require some consideration. At present, the most important issue is whether the operating system is patched when the customer first gets it, and subsequently. The UK House of Lords, for example, suggested mandatory ‘best-before’ dates on PCs, as these often sit in the supply chain for months and, once connected to the Internet, can be infected before the users even have time to connect to Microsoft to patch them up to date. Clearly, in such a case, the liability should fall on the shop rather than on the software vendor. Another solution would be to supply each PC with an up-to-date CD of patches; another might be to apply patches from a memory stick in the shop; yet another might be to redesign the software so that the machine would not connect to any other online service until it had visited the patching service and successfully applied an update. Regulation should seek to enforce the principle of security by default rather than engineer the details, which should be left to market players and forces. And we are careful to specify ‘all network-connected equipment’ rather than just PCs; if we see more and more consumer electronic devices online, but without mechanisms for vulnerabilities to be patched, then in due course they’ll be exploited.

‘Secure by Default’ isn’t just limited to patching. There are issues with active content (ActiveX, Visual Basic and JavaScript), which will no doubt change over time. Another issue is the provision of unneeded services. A vendor may bundle a web server on consumer PCs and printers, just to save installation costs – the idea being to install the same software on every machine and activate only those features that the customer pays for. However, if the unneeded services listen to the Internet, and let a piece of equipment get infected, then the liability must fall on the vendor.

The nature of certification also falls to be considered. One of the stakeholders expressed concern at the likely costs if all consumer electronics required Common Criteria



Source: Symantec ISTR vol. 9-12

Figure 12: Patch-development times for different operating systems

certification to EAL4; our view is that it would be quite sufficient for vendors to self-certify. However, the vendor should be liable if the certification later turns out to have been erroneous. Thus if a brand of TV set is widely compromised and becomes used for hosting phishing and pornography sites, the ISPs who paid penalty charges for providing network connectivity to these TV sets should be able to sue the TV vendor. Whether it was in fact the TV vendor's fault for having certified a TV as secure when it wasn't, or the distributor's fault for not patching it in time, is a matter for the court to determine in any particular case. (We expect though that once one or two landmark cases have been decided, the industry will rapidly adapt to a new liability system.)

In this way the Commission can start to move to a more incentive-compatible regime, by relentlessly reallocating slices of liability in response to specific market failures. It is also reasonable to make end-users liable for infections if they turn off automated patching or otherwise undermine the secure defaults provided by vendors. A useful analogy is that it's the car maker's responsibility to provide seat belts, and the motorist's responsibility to use them.

The next question is what other liability transfers should be made initially. The most important matters at the present time have to do with other aspects of patching – at which we must now look in greater detail.

6.5 Patching

Patching is an unfortunate but essential tool in managing the security of information systems. Patching suffers from two types of externalities. First, it is up to the software developer to create patches, but the adverse effects of a slow release are felt by consumers and the online community generally, rather than the companies directly involved. Second, the deployment of patches is costly, especially for large organisations. As discussed in the previous section, the publication of a patch often reveals the vulnerability to attackers, and

Vulnerability ID	Patch	Public exploit	Exploit appeared	Black market ad
Patch before exploit				
CVE-2007-3296	+2	N/A	+26	+29
CVE-2007-4105	+0	+62	+18	+52
MS07-004	+0	+7	+17	+13
MS07-009	+112	+153	+155	N/A
MS07-020	+0	N/A	+158	+105
MS07-027	+0	+2	+16	+26
MS07-035	+0	N/A	+29	+26
MS07-045	-1	N/A	+18	+18
Median (patch 1st)	+0	+34.5	+22	+26
Exploit before patch				
CVE-2007-3148	N/A	+0	+2	N/A
CVE-2007-4748	N/A	+12	+0	+11
CVE-2007-4816	+13	N/A	-1	+1
CVE-2007-5017	N/A	+0	+7	N/A
CVE-2007-5064	N/A	+20	+0	+15
MS07-017	+6	+11	+2	+13
MS07-033	+90	+0	+115	+91
Median (exploit 1st)	+13	+0	+2	+13

Source: Zhuge et al. [139]

Table 4: Time (in days) after public disclosure of vulnerabilities before a patch is issued and an exploit is published. The table also indicates when an exploit appears on Chinese websites and is advertised on the underground economy.

then the unpatched, compromised machines are used to harm others; so the local benefits of patching may be less than the local costs, even when the global benefits greatly exceed the costs.

6.5.1 Challenge 1: Speeding up patch development

The lag between vulnerability discovery and patch deployment is critical. During this period, consumers are vulnerable to exploits and have no recourse to protect themselves. So minimising this so-called ‘window of exposure’ is important. But software vendors are often slow in deploying patches, and there is great variation in the patch-development times exhibited by different vendors. Figure 12 plots the patch-development times for several operating system vendors during the past two years. Microsoft and Red Hat are fastest, Sun and HP are slowest by far, and Apple is in the middle. Consumer-oriented OSs tend to patch faster, perhaps because there is greater consumer demand and awareness for security updates.

It is also important to understand the relationship between the availability of patches, the creation of exploits, and the exploits’ use in the underground economy. Table 6.5.1 indicates the time difference between the publication of vulnerabilities and the appearance of patches and exploits for vulnerabilities exploited by Chinese websites in 2007 [139]. The

top portion of the table shows vulnerabilities where patches are released before exploits are observed, while the bottom portion lists vulnerabilities where exploits appeared in the wild before patches were available.

Nearly half of the vulnerabilities in Table 6.5.1 were actively exploited in the wild before a patch was disclosed. Notably, the median time lag between the vulnerability being disclosed and it appearing in the wild is just two days, while patches took nearly two weeks to be published (if they were released at all). This suggests that there is scope for speeding up patch dissemination.

Option 1: Responsible vulnerability disclosure Vulnerability disclosure is often what triggers the development and deployment of patches. Yet the process by which the vulnerability is disclosed can affect the time vendors take to release patches. Some security researchers advocate full and immediate disclosure: publishing details (including potentially exploit code) on the Bugtraq mailing list [122]. While undoubtedly prompting the vendors to publish a patch, full and immediate disclosure has the unfortunate side effect of leaving consumers immediately vulnerable. Vendors, for their part, typically prefer that vulnerabilities never be disclosed. However, some vulnerabilities might go undiscovered by the vendor even when they're being exploited by miscreants, and non-disclosure creates a culture in which vendors turn a blind eye.

A more balanced alternative is responsible disclosure as pioneered by CERT/CC in the US. CERT/CC notifies vendors to give them time to develop a patch before disclosing the vulnerability publicly. When the vulnerability is finally disclosed, no exploit code is provided.

Empirical analysis comparing the patch-development times for vulnerabilities reported to Bugtraq and to CERT/CC revealed that CERT/CC's policy of responsible disclosure led to *faster* patch-development times than Bugtraq's full disclosure policy [9]. This is because CERT/CC has developed a more constructive relationship with software vendors, working with them to fix vulnerabilities. The researchers also found that early disclosure, via CERT/CC or Bugtraq, does speed up patch-development time.

Option 2: Vendor liability for unpatched software Another option is to assign liability for vulnerabilities to the software vendor until a patch is made available and consumer has reasonable chance to update. This could encourage faster patching.

Cavusoğlu et al. compare liability and cost-sharing as mechanisms for incentivising vendors to work harder at patching their software [22]. It turns out that liability helps where vendors release less often than optimally.

Option 3: Fixed penalty for slow patchers Since liability has been fiercely (and so far successfully) resisted by software vendors, it is worth considering alternatives that could also speed up patch deployment. Vendors slow to issue patches could be charged a fixed penalty. Given that some operating system vendors are much slower to release patches than others (Figure 12), a fixed penalty may be quite effective at improving the overall speed of laggards.

One drawback of fixed penalties based on a single time threshold, however, is that most vendors already prioritise their patch development to push out fixes to the most

severe vulnerabilities fastest. Introducing a time-based penalty may draw resources away from developing critical patches in favour of less-important ones near the deadline.

Recommendation 6: We recommend that the EU adopt a combination of early responsible vulnerability disclosure and vendor liability for unpatched software to speed the patch-development cycle.

6.5.2 Challenge 2: Increasing patch uptake

While quantitative measurements are difficult to obtain, the view among security professionals is that patches are available for the majority of exploits used by attackers. Over half of the exploits in Table 6.5.1 appeared on Chinese websites after a patch was made available. Because these exploits are being advertised well after the patch was available, this provides evidence that attackers target unpatched machines. Judging from the median values (22-day lag for patched vulnerabilities versus 2-day lag for zero-day exploits), whenever patches are published before exploits, attackers are less rushed to develop exploits since the target will be unpatched systems, and presumably, they will continue to be unpatched for a long time.

So why do some users remain unpatched? While most operating systems offer automatic patching, many third-party applications like web browser add-ons do not. Some perfectly rational users (especially at the enterprise level) choose not to patch immediately because of reliability and system stability concerns. Quantitative analysis of security patch deployment reveals that pioneers end up discovering problems with patches that cause their systems to break [12]. Typically, waiting ten to thirty days best serves a business's own interests.

Option 1: Free security patches kept separate from feature updates Vendors must make patching easier and less of a nuisance for consumers. One simple way of doing this is to decouple security patches from feature updates. Users may not want to add the latest features to a program for a variety of reasons. Feature updates could disrupt customisation, slow down performance, or add undesirable features (e.g., DRM). Even though most feature updates are beneficial, the few disruptive updates could turn off users to patching, even when it is in their interest to do so.

Microsoft's Windows Genuine Advantage (WGA) program is an anti-piracy tool that users are required to install before downloading updates. WGA provides a useful example of how meddling with the update process can turn off users to patching. Rather than treating validation as a one-off process, in its initial design WGA connected to Microsoft following every boot-up. This triggered outrage from privacy advocates, to which Microsoft eventually yielded. One positive aspect of WGA, by contrast, is that it allows even pirated software to be eligible for security patches. Other companies should do the same.

Microsoft again violated the trust of many users when it emerged that Windows Update automatically installed new updates even when users had explicitly asked for approval first [87]. Software companies should make the updating process as transparent as possible, given the importance of patching.

Option 2: Vendor liability for software without automated patching Some types of software do not offer automated patching. This introduces an unacceptable burden on users. Vendors who do not provide automated patches could be held liable. This could be implemented as part of the ‘safe default’ approach to liability discussed in Section 6.

Option 3: Vendor-firm cost-sharing Installing patches at the enterprise can be expensive, imposing significant IT labour costs for verification and troubleshooting. At the same time, firms may not see the benefit of patching, particularly when attacks target third parties. One solution is for the software vendor to subsidise the costs of patch installation at the vendor. This could be negotiated between the vendor and firm, so it is unclear whether regulation is needed.

Recommendation 7: We recommend security patches be offered for free, and that patches be kept separate from feature updates.

6.6 Consumer policy

Where consumers are involved one may need more protection. Competition is relevant here too: consumers are in a weak position vis-à-vis competing vendors of products where there is an ‘industry position’ of disclaiming liability for defects (as with cars two generations ago, or software and online services today), yet they are in an even weaker position facing a monopoly supplier such as Microsoft. In both cases, they are faced with shrink-wrap or click-wrap licenses that impose contract terms on them, on a take-it-or-leave-it basis.

Shrink-wrap licenses are thought by legal scholars to be defective: they attempt to impose terms after the purchase of a product, so in effect you’re not buying the product but an option to enter into a license agreement provided you haven’t done things you already in fact have done. Lawyers argue that this is like a hotel pinning up terms and conditions inside the wardrobe door – it’s too late. However as firms move to software download and click-wrap, this issue may become moot. In any case, citizens need consumer protections that are properly engineered and fit for purpose, rather than just relying on the side-effects of a transient technology for questionable protection.

6.6.1 Fair contract terms

The main applicable law in the EU is based on the Unfair Contract Terms Directive [50], which makes a consumer contract term unfair ‘if, contrary to the requirement of good faith, it causes a significant imbalance in the parties’ rights and obligations arising under the contract, to the detriment of the consumer’. This is widely flouted by the software industry. For example, Article 5 requires that ‘terms must always be drafted in plain, intelligible language’; yet in practice, end-user license agreements (EULAs) are written in dense legalese and made difficult to access; a large amount of text may appear via a small window, so that the user has to scroll down dozens or even hundreds of times to read it. Article 7 further requires Member States to ensure that ‘adequate and effective means

exist to prevent the continued use of unfair terms in contracts concluded with consumers by sellers or suppliers’.

Some Member States have even stricter laws, the UK being an example [27]; and in some circumstances, unfair-contracts law has also been used by firms or public bodies against suppliers. A well-known case is *St Albans District Council vs ICL*. ICL sold the council software containing bugs that caused financial losses; the council sued, and the court found not only that the software was not fit for purpose, but that the Unfair Contract Terms Act applied because the council signed the unmodified Standard Terms and Conditions provided by ICL [127].

There remain many areas, though, in which unfair terms for both software and services persist, despite the fact that in theory they could be challenged in the courts. Again, banking provides an example: Bohm, Brown and Gladman analyse how, when banks rushed to set up online banking services during the dotcom boom, many of them changed their terms and conditions so that customers who accepted passwords for use in electronic banking also accepted liability for all transactions where the bank claimed that their password had been used [13]. The liability for fraud and security failure in online banking was thus transferred (at least on paper) to the customer.

There is significant variation across Member States in how complaints about fraudulent electronic banking transactions are handled. In both the UK and Germany, banks have transferred liability, generally to customers where a PIN or password is used, and to the merchant for signature-based or online transactions. However, the practical consequences for customers differ; in the UK, court rules that the loser in a civil matter must pay the winner’s costs make it impractical for most people to sue, while in Germany a bank can recover only very limited costs from a customer who sues it and loses; thus in practice German bank customers are better protected. The UK has a ‘Financial Ombudsman Service’ that provides alternative dispute resolution between banks and customers; this service is without cost to the customer, being paid for by the banking industry, but has been accused of partiality towards the banks and is currently the subject of a review by Lord Hunt. In the Netherlands, the banks claim to always refund defrauded customers but have resisted any actual legal liability. Ireland is also important, as the seat in Europe of PayPal; PayPal, like the Dutch banks, claims to have always made good every customer who has been the victim of fraud, and yet their terms and conditions specify that disputes should be resolved by reference to the UK Financial Services Ombudsman. By way of comparison, the US Regulation E, which governs electronic banking, places the onus of proof squarely on the bank – which as the operator of the electronic payment system is the only party in a position to really affect the fraud rate. This is not merely because it designs and maintains the payment system itself, but because it has access to deep and wide information about the patterns of fraud across many merchants and customers.

The question of varying fraud liability and dispute resolution procedures has been raised from time to time, and so far has been avoided by legislators (most recently when the Payment Services Directive was being negotiated from 2002–5 [62]). We believe the time has come for the Commission to tackle this issue.

Recommendation 8: The European Union should harmonise procedures for the resolution of disputes between customers and payment service providers over electronic transactions.

6.6.2 Protection against abusive practices

Some companies use deceptive marketing techniques that break various EU laws. Spyware programs ‘monitor user activities, and transmit user information to remote servers and/or show targeted advertisements’ [39]. Spyware is bad for several reasons. First, it often employs deceptive installation practices: piggy-backing on installations of other programs, exploiting security holes, or using unsolicited ActiveX pop-ups while browsing web sites [37]. These installation strategies violate the Unfair Contract Terms Directive. In almost all cases, the installation will be done without valid, free consent, so spyware users violate the Data Protection Directive and the E-Privacy Directive [58]. As if that weren’t enough, spyware programs are often made deliberately hard to uninstall.

Once installed, spyware collects extensive data on user behavior without user consent, in violation of data protection legislation. Spyware effectively hijacks the advertising channel for web browsing. Many merchant websites pay a commission to affiliate websites whenever a user follows a link from the affiliate website to the merchant. Spyware intercepts this process to claim the commission for the spyware vendor. So spyware is a problem not only for consumers, but also SMEs running websites that rely on affiliate revenue.

Dealing with spyware through regulation is difficult, since most spyware companies are based outside the EU (typically in the US). US regulators are trying to rein in the excesses of these companies [134], but looser laws mean that they are allowed to carry out dodgy practices that are forbidden in the EU. Furthermore, there is evidence that the terms agreed between spyware vendors and US regulators are being flouted [40].

While directly regulating the practices of spyware vendors is difficult, effective sanctions are still possible by punishing the companies that advertise using spyware. In the 1960’s, a number of unlicensed ‘pirate’ radio stations aimed at UK consumers were launched from ships just outside the UK’s jurisdiction. The Marine Broadcasting Offences Act of 1967 made it illegal for anyone subject to UK law to operate or assist the stations. This immediately dried up advertising revenues, and the unlicensed stations were forced to fold. A similar strategy could undermine spyware, since many of the advertisers are large international companies that do business in the EU [38]. While advertisers might object that they could be framed by competitors, an examination of the resulting evidence should vindicate any false accusations.

Another abusive practice already the target of regulation is spam. The EU Directive on privacy and electronic communications [58] attempts to protect consumers from spam. For the most part, it prohibits sending any unsolicited messages to individuals, requiring their prior consent. However, there are two exemptions worth discussing.

The first exception comes from Article 13 paragraph 2. It allows for unsolicited communications provided the consumer has bought something from the company in the past and is given a clear opportunity to opt out of receiving the messages. The Commission struck a balance in setting this exception. It remains tractable for consumers to individually opt out of spam arising from previous transactions. Individually opting out of spam sent by many thousands of companies where no prior business relationship exists, by contrast, would cause undue burden. As such, we support this exemption.

A second exception arising from Article 13 paragraph 5, however, is more problematic. This paragraph states that protections only apply to ‘natural persons’, and leaves it up to Member States to decide whether to allow unsolicited communications to business.

Direct marketing lobbies argued that spamming businesses was essential to their trade. In practice, the business exemption has undermined the protections for consumers. It gives spammers a defence against all messages sent to ‘work’ domains. It also drives up costs for businesses, who must contend with spam sent from potentially millions of other businesses. Finally, it is also difficult (in practice impossible) to draw clear lines between ‘natural’ and ‘legal’ persons in this context: some businesses (one-man firms, barristers, partners in some organisations) are legally ‘natural’ persons, while email addresses of identifiable individuals in companies relate to ‘natural’ persons. So there is a strong case to abandon the distinction. Therefore, we recommend repealing Article 13 paragraph 5, the business exemption for spam.

Putting all these together:

Recommendation 9: We recommend that the European Commission prepare a proposal for a Directive establishing a coherent regime of proportionate and effective sanctions against abusive online marketers.

6.6.3 Consumer protection in general

The issues raised in this section on consumer policy are not limited to abusive marketing and unfair banking contracts. There are many more problems on the fringes of information security that warrant further study.

For example, as e-commerce becomes m-commerce, abusive practices in the telecomms industry are becoming increasingly relevant. These include *slamming* (changing a customer’s phone service provider without their consent) and *cramming* (dishonestly adding extra charges to a phone bill). For example, one of us was the victim on an attempt at cramming. On holiday in Barcelona, a phone was stolen when a bag was snatched, and the account was immediately cancelled. Several months later, the mobile service provider demanded payment (of a few tens of euros) for roaming charges recently incurred by that SIM in Spain. In all probability, the Spanish phone company was simply cramming a few charges on a number they’d seen previously, in the knowledge that they’d usually get away with it. It took substantial argument with the mobile service provider to get the charges dropped, requiring escalation to the chairman’s office. Mobile service providers find it easier to blame customers than to argue with business partners, and a recent trend is to sell customers ‘insurance’ to cover such disputed calls. This appears to be a clear regulatory (and policing) failure.

A second example comes from ‘identity theft’. This is actually a misnomer; Adam Shostack and Paul Syverson argue persuasively that identity theft is actually libel [125]. Fifteen years ago, if someone went to a bank, pretended to be you, borrowed money from them and vanished, then that was the offence of impersonation and it was the bank’s problem, not yours. In the USA and the UK in particular, banks have recently taken to claiming that it’s your identity that’s been stolen rather than their money, and that this somehow makes you liable. The situation does not yet appear to be as bad in other Member States (many of which do not yet have the UK/US culture of credit histories as ‘financial CVs’) but that is no reason for complacency (as the UK/USA culture is spread by the pressures of globalisation).

A bank should bear full liability for the consequences of mistaking an innocent person

for a third party and should not pass on false and defamatory information on that person to the credit-reference agencies. In theory, the data protection authorities could compel a bank or an agency to cease and desist from knowingly disseminating false and defamatory information about an individual, but in the UK at least the authorities have declined to do this. A further option, which is increasingly common in the USA, is credit locking: a citizen who does not want any more credit – for example, a middle-aged person who’s paid for their house and has enough credit cards – simply forbids the credit-reference agencies to give any information on them to anyone. However, in the UK the agencies charge a significant sum for this service. This appears also to be a regulatory failure.

Our third, and perhaps most important, example concerns the foundation of the Single Market itself. The European Union has long been more than a ‘Zollverein’ and it is a long-established principle that citizens can buy goods anywhere in the Union. (As one US lawyer put it to us, ‘You’ve elevated grey-market trading into a fundamental human right!’) It is rational for firms to charge discriminatory prices; as people earn more in London than in Sofia, a clothing vendor will naturally charge more for trousers there. But this is unpopular and it has long been policy that anyone may buy trousers in Sofia, put them on a truck, take them to London and sell them. Now the value of physical goods is often tied up with intellectual property, such as a trade mark, and the Union has had to develop a doctrine of first-sale exhaustion to deal with that. The challenge now is that goods are increasingly bundled with online services, which may be priced differently in different Member States, or even unavailable in some of them. The bundling of goods and services is an area of significant complexity in EU law. Sometimes the problem is solved when a market becomes more competitive (as with personal video recorders over the past few years) but sometimes the market segmentation persists.

The relationship between the segmentation of online service markets and information security is complex. For example, during the 1990s, Sky TV stopped broadcasting Star Trek in Germany, and this led many German students to investigate ways of breaking pay-TV security. This led in turn to the discovery of many vulnerabilities in the smartcards of the time, and to several rounds of attack-defence coevolution in hardware tamper-resistance [5]. And, as already noted, national laws already segment markets: Flickr provides a more restricted service to customers in Germany out of (probably misplaced) concerns about obscenity. Sometimes market segmentation in B2B transactions has an effect on consumers; for example, citizens in one country can find it hard to open a bank account in another because of the way in which credit-reference services are bundled and sold to banks. This in turn reduces consumers’ ability to exert pressure on banks in countries where online banking service is less competitive by switching their business elsewhere.

The 2006 Services Directive takes some welcome first steps towards harmonising the market for services [61], seeking to remove legal and administrative barriers in some fields (such as hotels, car hire, construction, advertising services and architects) while unfortunately excluding others (including broadcasting, postal services, audiovisual services, temporary employment agencies, gambling and healthcare). This Directive focuses on removing the many protectionist measures erected over the centuries by Member States to cosset domestic service providers, and rightly so. In our view however there is another aspect, namely the deliberate use of differential service provision as a tool by marketers, both as a means of discriminatory pricing and in order to undermine consumer rights.

Single-market service provision is very much broader than the scope of this report; it encompasses issues from extended-warranty insurance through frequent-flyer programs. Like the liability for defects in software – and in services – it’s such a large topic that it will have to be tackled a slice at a time, and by many stakeholders in the Commission. We encourage ENISA to become involved in this policy process so that the security (and in broader terms the dependability and safety) aspects of policy are properly considered along with the straightforward consumer-protection questions.

Finally, the issue of universal access to the Internet, to which we referred in the discussion on Recommendation 4, may also benefit from action under the heading of consumer rights. If all the ISPs in a country align their terms and conditions so that they can disconnect any customer for no reason, this should be contrary to public policy on a number of grounds, including free speech and the avoidance of discrimination. For example, legal action was taken by the Scientologists to suppress material made available via the Finnish remailer `anon.penet.fi` and the Dutch ISP XS4all [5]; and one of us (Anderson) was once the target of harassment by animal rights activists by virtue of his being a member of his university’s governing body. Even those citizens who are unpopular with some vocal lobby group must have the right to Internet connectivity. The Commission should give thought as to how this right is to be defended.

Recommendation 10: ENISA should conduct research, coordinated with other affected stakeholders and the European Commission, to study what changes are needed to consumer-protection law as commerce moves online.

7 Dealing with the lack of diversity

Diversity, as a security property, can be described as the absence of single points of failure. We distinguish physical diversity from logical diversity. Physical diversity deals with geographical distribution of redundant infrastructure components and the routes of network fibre connecting them, whereas logical diversity means that distributed systems do not share common design or implementation flaws. While physical diversity has been an issue for long, the importance of logical diversity increases with the degree of system interconnectedness and the ability of strategic attackers to exploit vulnerabilities remotely (thus thwarting efforts of physical diversity). A lack of diversity implies risk concentration which negatively affects insurability and thus an economy's ability to deal with cyber risks. Unfortunately, free markets often work against diversity, which explains calls for government intervention.

7.1 Promoting logical diversity

For logical diversity to happen, alternatives must be widely available and adoption well-balanced. In practice, this has rarely occurred due to the structure of the IT market: fast technology cycles, positive network externalities and high switching costs between technologies tend to yield dominant incumbents and fading competition [124]. Nonetheless, there are steps governments can take to improve, or at least not hinder, the prospects for diversity.

Option 1: Promoting open standards to facilitate market entry A policy to foster diversity must first ensure the availability of viable alternatives. One option is to promote open standards to facilitate market entry. Open standards are no panacea, but they allow competitors to develop interoperable software and crack customer lock-in, one strong force which otherwise keeps customers in the incumbent's claws.

Notably, open standards are also on the agenda of the European Commission's Interoperable Delivery of European eGovernment Services to Public Administrations, Businesses and Citizens (IDABC) initiative¹⁵, albeit to ensure interoperability and competition rather than to improve security. It would be useful for ENISA to liaise with IDABC so that whenever diversity has security implications this is brought to the fore. This effort could complement ENISA's activity on specific security standards.¹⁶

However, promoting open standards is not the same as promoting diversity. Microsoft has heavily promoted its rival standard [36] to the Open Document Format (ODF) [84]. By using containers for proprietary data, compatible alternative implementations may be squashed, yielding the same dominant outcome as already exists.

Even successful open standards often do not lead to diversity. Most applications supporting the Portable Network Graphics (PNG) format across platforms rely on the same reference implementation library *libpng*¹⁷ for image processing. As a result, vulnerabilities in one library (of which there are many: 17 vulnerabilities for *libpng*, including 5 critical,

¹⁵see <http://europa.eu.int/idabc/>

¹⁶see the ENISA/ITU ICT Security Standards Database launched in June 2007 <http://www.itu.int/ITU-T/security/main.table.aspx>

¹⁷<http://www.libpng.org/>

according to the National Vulnerability Database¹⁸) can lead to multi-platform exploits. The *libpng* library is but one example of how hidden homogeneity at the lower levels can wreak havoc even when applications and systems platforms appear superficially diverse.

Option 2: Promoting diversity in the procurement process and e-Government

Consumers and firms are understandably short-sighted when selecting a software product. The positive network externalities of user adoption mean that they are likely to ignore any increase in correlated risk. Governments, however, need not be so myopic. They can encourage the adoption of rival technologies during public procurement. Unfortunately, they often pursue policies detrimental to diversity.

In 2004 the European Commission examined public procurement practices for IT equipment in several Member States and found that the specifications for the requested processor architecture favoured Intel products. This directly strengthened the dominant platform [59]. Although only France, the Netherlands, Finland and Sweden were explicitly mentioned, other countries including Germany and Ireland changed their procurement rules in reaction to the EC call.

Another example comes from Germany, where most businesses are required to submit tax statements electronically. However, the ELSTER software used to submit annual trade tax and VAT statements is only fully compatible with the Windows platform. Small businesses considering a migration to alternative platforms must know that they can submit their forms three years later. At present, this is not certain as the software is currently revised every year to reflect the latest changes in the tax code.

When citizens interact with their government online, they are often required to use Microsoft Office formats only. Governments should provide a better example by offering documents in several formats.

There have been several positive examples of governments choosing less dominant software platforms, albeit for cost-saving reasons¹⁹. After a heated debate, the German Bundestag, the lower house of the federal parliament, decided in 2002 to replace its server infrastructure in large parts with one that runs a Linux operating system and uses OpenLDAP, an open standard for directory services, to connect with several thousands of Windows desktop computers [75]. The city of Munich went a step further by installing Linux on 14,000 desktop PCs of the city administration which run 1,100 different applications altogether.²⁰ Other cities and countries have followed, from the city of Vienna to the French government, which spent 11 % of public IT expenditure on open source software in 2007 [126] and ran OpenOffice on 400,000 workstations [48].

Option 3: Advise competition authorities when lack of diversity presents a security issue There are limits to the impact governments can have through public procurement policies alone. Regulatory responses may occasionally be required if the

¹⁸<http://nvd.nist.gov>

¹⁹While the following examples all involve open source software, this is unintentional and not relevant to the case for diversity. Rather, it is a reflection of the fact that few commercial alternatives exist at present. A thorough economic analysis of government funding of open source software is orthogonal to most NIS aspects and therefore beyond the scope of this report. We refer the reader to the relevant literature instead [120].

²⁰see <http://www.muenchen.de/Rathaus/dir/linux/english/147197/index.html>

security threat is high enough. As already mentioned, diversity is often rightly viewed as a competition and consumer issue. So it makes sense for ENISA to take an active role in advising the competition and consumer regulators whenever diversity presents a security threat.

As mentioned earlier, Cisco used to have a very dominant market position in the routers deployed in the Internet backbone. A vulnerability in Cisco routers [137] was disclosed that could remove a significant portion of the Internet backbone if a flash worm was disseminated. Hence, the lack of diversity among routers used to be a critical concern. However, the market for backbone routers has balanced recently, given competition from Juniper and other companies. The market for mobile-phone software similarly used to be dominated by Symbian, but that has also corrected itself somewhat thanks to challenges by Apple, Google, Microsoft and others. Finally, the market for web browsers is now more competitive following years of dominance by Internet Explorer.

In each of these cases, market forces have eventually helped to mitigate the lack of diversity in products. However, some dominant products have resisted repeated action by the competition authorities – Windows comes to mind – and regulators need to be aware of security threats that follow from lack of diversity, in addition to the competitive threats. ENISA, with input from technical experts, could take this role.

Recommendation 11: We recommend that ENISA should advise the competition authorities whenever diversity has security implications.

7.2 Promoting physical diversity in CNI

The critical national infrastructure (CNI) comprises the systems and services that underpin the economic, social and political structures of a nation. It is usual to include communications in general, and – increasingly since the mid-1990s – the Internet in particular as one part of the CNI. Pitcom, a UK parliamentary group, has published a useful overview aimed at legislators [113]. They pick out two specific threats to the Internet – ‘hacking’ and damage to ‘choke points’, then go on to show how an Internet failure would damage other parts of the CNI such as Finance, Food and Health.

This interconnection between parts of the CNI is increasingly common; if a high voltage power line fails the engineers who go to fix it will keep in touch with their base by mobile telephone. But the mobile telephones depend on the public power supply to keep base stations operating. Self-contained ‘satellite phones’ would solve this problem, but they are expensive to own and operate, so cost-saving measures may mean that insufficient numbers are purchased.

7.2.1 Common mode failures and single points of failure

In principle, ‘choke points’ are avoided by communications network designers, who call them ‘single-points of failure’ and introduce redundant components to design them out. However, they may be beyond an individual network’s control or the failure may be beyond their imagining. The Buncefield oil refinery explosion in December 2005 severely damaged a Northgate Information Solutions building, taking out systems for over 200 different customers, including payroll systems for over 180 clients and patient administra-

tion systems for hospitals as far away as Cambridge and Great Yarmouth. The damage from ‘the largest explosion in peacetime Europe’ was so extensive that onsite backup systems were also obliterated and offsite facilities had to be brought into use, with downtimes measured in days. Designers are regularly caught out by common-mode failures, whether it be by putting backup systems in the other World Trade Center tower [34], purchasing communications links from different companies that end up going over the same bridge that is washed away in a flood, or having vandals pour petrol down into underground cable ducts carrying many disparate cables and then setting them on fire [97].

Efforts are being made to improve information about common-mode failures, and customers are increasingly insisting on knowing where fibre actually runs when they purchase telecomms circuits. Other lessons are being learnt from 9/11, in particular that systems switched to backup power, but that refuelling arrangements used a small number of companies – who could well have been overstretched, but in the event they couldn’t get permission to enter lower Manhattan anyway. Although there were schemes for getting priority access, the only companies involved were those that were in existence in the 50’s and 60’s when planners were considering nuclear war. The modern ‘dot-com’ companies were completely outside of these systems. In London Docklands there are now regular planning meetings between police, local authorities, data centre operators, Internet companies, and so on. In the event of an incident, there may still be difficulties in accessing the Docklands area while it remains a ‘crime scene’, but at least the police have been educated into understanding why that access might be necessary.

7.2.2 Internet exchange points

A major concern about single points of failure for the Internet is the growth of Internet Exchange Points (IXPs) such as LINX in London, AMSIX in Amsterdam, DECIX in Frankfurt etc, and the way in which there are tendencies towards one IXP becoming significantly larger than its rivals.

ISPs need to be able to provide their customers with connectivity to the whole of the rest of the Internet. They do this by purchasing ‘transit’ from a major networking company, paying for their traffic on a volume basis. To reduce their costs ISPs will attempt to negotiate ‘private peering’ arrangements with other ISPs, where traffic is exchanged ‘settlement free’. This traffic will not be for ‘all possible routes’, but only for the parts of the Internet operated by the other ISP. The largest ‘backbone’ networks (usually called Tier 1 networks) do not purchase transit from anyone, but operate solely on a peering basis. In the past there were only about 5 Tier 1 networks, but there are probably 9 at present, with another 20 or so ‘Tier 2’ networks that have peering-only arrangements in large geographical regions, but use a Tier 1 for remote locations.

ISPs often use IXPs to reduce the costs of peering. One of the ISP’s routers is housed at the exchange point and ‘public peering’ traffic (and possibly transit traffic as well) is exchanged with other ISPs over the fabric of the IXP (multiple high speed Ethernet rings at larger IXPs, a switch backplane at the smallest). An ISP with large numbers of customers (or ‘eyeballs’, viz: a data sink) will find it easy to arrange peering with an ISP with a large number of content providers (a data source) because it will be in both their interests to avoid paying for transit. Thus, for example, companies such as Google, Akamai and the BBC will generally peer with anyone.

Although the trend is towards specialisation, many ISPs, particularly those who have

been in business longest will be a hybrid of eyeballs and content. In general, these hybrids will be able to set up peering arrangements with other hybrid ISPs of the same size. Companies generally refuse to peer with ISPs that are a lot smaller than themselves, taking the view that ‘they ought to be a transit customer’.²¹ Similarly, ISPs that mainly exist to haul traffic (often described as NSPs – Network Service Providers) will refuse to peer with ISPs that they believe should be purchasing transit.

The value of joining an IXP can clearly be seen to increase as more ISPs join, so that there is an obvious economic pressure towards winner-take-all scenarios where one IXP is much larger than its local rivals. In 11 EU countries there is just one IXP, in almost all the others the largest IXP is 4 or more times the size²² of the next largest – the exceptions being Estonia, Spain, Belgium, and Poland (in each of which there are 2 roughly equal size IXPs, an unstable non-monopoly equilibrium) and France which, for complex historical reasons, is much more fragmented with 5 similar sized exchanges. Table 13 in Appendix B gives a complete list of European IXPs and their size. The smaller exchanges, where they exist, are viable either because their members are of a different size to those at the main IXP (so they already have most of the peering they would get at the main IXP) or because of complex historical situations where ISPs found it too difficult (for reasons of price, or perceived non-neutrality) to join the main IXP.

Unfortunately, the economic pressures towards a dominant IXP could lead to single-points of failure when there is a problem with the IXP itself. The largest IXPs deal with this through diversity within the IXP itself. For example, LINX operates in multiple buildings in London Docklands with two physically separate peering LANs from two different vendors, so that there is little chance of a common-mode failure. AMSIX in Amsterdam has an entire redundant fail-over system. However, not all IXPs have taken such steps, mainly because of the expense.

For larger ISPs there is no problem; they will be connected to IXPs in multiple countries, so if AMSIX fails they can exchange traffic at LINX and vice versa. However, smaller ISPs cannot afford international links, so an IXP failure will increase their costs (as they have to use transit for all of their traffic). It may even cause partial or complete failure for their customers if the transit link cannot handle the traffic, or if their transit traffic goes via the IXP as well.

There has been some regulatory interference with these arrangements in that the Access and Interconnection Directive [54] makes it unlawful for one network to refuse to interconnect with another, although this connection will be made at commercial rates. This measure was mainly aimed at telephone networks, where some of the newer ‘alternative’ telcos, particularly in the mobile market, were finding it hard to persuade the incumbents to interconnect.

The Directive has made little or no difference to Internet transit provision, where a highly competitive market is keeping prices low. Indeed as new higher-capacity links become available, and prices drop, it is often worthwhile to cancel existing contracts that still have months to run, pay penalty clauses, and make a new contract with another

²¹Historically, IXPs were set up by the hybrid ISPs, who for some time resisted the attempts of content providers to join the IXP (because the hybrids hoped to sell transit), but most European IXPs have now swept away the complicated joining conditions of the past and let most anyone become a member.

²²Size is being measured here in terms of number of customers at the IXP, rather than the volume of traffic exchanged.

provider.

However, the Directive, and other legislation on fair competition, has occasionally had a perverse effect on peering arrangements, and thereby reduced redundancy for Internet traffic flows. The largest ISPs will, as explained above, achieve resilience by connecting to IXPs in other countries. Historically, they would then peer with pretty much every other ISP at that exchange because they didn't expect those other ISPs to become customers – and the peering would save them money on transit. However, they would not peer with small IXPs in their home country – because they were (or 'ought to be') customers. But, to avoid any risk of being caught by non-discrimination rules, these large ISPs have sometimes refused to peer with small ISPs at foreign IXPs, apparently because they may be forced by the regulator to provide free peering at their home IXP.²³

7.2.3 Hacking the critical national infrastructure

There is widespread concern about 'hacking' damaging the CNI, although there are only a handful of known cases, and few of the events involve malicious outsiders specifically targeting the CNI system itself. In the USA, attention has focussed on SCADA devices connected to the Internet without due consideration to making them secure. The economics have been different in Europe, so there are rather fewer such systems. Nevertheless in December 2004 the EU agreed a European Programme for Critical Infrastructure Protection (EPCIP) and a Critical Infrastructure Warning Information Network (CIWIN), and in November 2005 the Commission adopted a green paper on a European Programme for Critical Infrastructure Protection (COM (2005) 576 final) [44].

The other, entirely European 'hacking' event of note is the April/May 2007 denial-of-service attack on Estonia. This event created widespread alarm, claims of involvement by the Russian Government, and claims that the first 'cyberwar' was taking place. However, careful measurements showed that the attacks were only of the order of 90 Mbit/s which is really quite small (Japanese consumers can purchase 100 Mbit/s links for approximately USD 50 per month). The real problem was that Estonia had a fairly low-bandwidth infrastructure, and a lack of experience in dealing with DDoS attacks, so significant problems arose from a relatively small attack. Lesk [91] estimated at the time that if botnets had been rented specially for the purpose, each of the attacks on Estonia would have cost only a few thousand dollars – and since then a 20-year-old ethnic Russian has been convicted of the attacks and fined 17,500 kroons (EUR 1,100). Estonian officials now admit they have no other suspects [66].

7.2.4 Policy options

Critical National Infrastructure is now understood to be a multi-national issue, and we have already noted the initiatives made by the Commission in this area. However, there has been no formal follow-up to their Green Paper [44].

²³Unfortunately, peering arrangements are seldom public, so it is not possible to provide a citation for this claim. Indeed, some industry experts suggest that it may not be the case, and the explanation for the observed behaviour is just that the large IXPs are too inefficient to get around to setting up all the peering that would be economically efficient. However, examination of the nature of the firms involved in the rather small number of occasions where long-standing peering arrangements have been terminated gives some credence to our explanation.

One of the key difficulties in this area is that it is dominated by secrecy (CNI companies do not wish to discuss how they might be vulnerable) and by limited understanding of the real world: for example the COCOMBINE project in Framework 6 examined IXPs as a part of its work. However, it failed to understand why peering does or does not take place between particular ISPs, and merely attempted to find spatial patterns, with limited success [79, 72, 80].

We earlier remarked that when AMSIX has a problem the traffic is expected to go via LINX. This is based on observations of a handful of historic events. However, whether this remains the case today, or whether the traffic might traverse a more minor IXP instead (causing it in turn to fail) is clearly of significant interest to disaster planners.

Hence the most obvious policy option to adopt is that of encouraging information sharing and more, and better informed, research into the actual issues. Scaremongering about ‘cyberwar’ has proved effective at unlocking research coffers at the US Department of Homeland Security, but without more information about specifically European issues, it is hard to even scaremonger effectively.

The other obvious policy option is of sharing and promoting Best Practice. For example, in the UK the major IXP is LINX, and it has deliberately chosen to run two co-located but physically distinct Ethernet peering rings so as to provide significant resilience. When it found that ISPs were not bothering to connect to the second ring it changed its charging structures to make it cheaper to connect to the second ring rather than purchase more bandwidth on the first. It also monitors the extent to which members connect in one main building (Telehouse) rather than the other six nearby locations at which it has a presence. Many other European IXPs do not have this level of diversity.

The other option is of course regulation. As we have already noted, well-meaning regulation on interconnection may have had the perverse effect of reducing resilience, and increasing costs. Without significantly better understanding of the issues, this is not an option that can be recommended. In our view, the appropriate level of compulsion is given by the following recommendation.

Recommendation 12: We recommend that ENISA sponsor research to better understand the effects of IXP failures. We also recommend they work with telecomms regulators to insist on best practice in IXP peering resilience.

8 Fragmentation of legislation and law enforcement

8.1 Criminal law

To a first approximation, existing legal frameworks have had no difficulty in dealing with the Internet. Whether criminals use letters, telegrams, telephones or the Internet, fraud is fraud, extortion is extortion, and death threats are death threats. The mantra ‘if it’s illegal offline it’s illegal online’ has been effective at calming those who see new threats to civilised life in the new medium, and it has only been necessary to construct a handful of novel offences that can only be committed in cyberspace. The first such attempt at setting out these offences was the UK’s Computer Misuse Act 1990 [132]:

- Existing notions of trespass were inadequate for criminalizing computer hacking, so specific offences for unauthorised access to computers were put in place.
- Offences were constructed for the creation and distribution of computer ‘viruses’.

Since 1990, with the advent of the Internet as a mass medium, this list has been extended with:

- Offences for denial of service attacks (where the network itself is the target rather than individual machines per se).
- Forbidding collections of hacking tools and passwords (where these collections are possessed ‘without right’).

However, the cross-jurisdictional nature of cyberspace has meant that many criminals commit their offences in another country (often many other countries) and this leads to difficulties in ensuring that they have committed an offence in the country in which they reside. This is not a new problem. Brenner [18] notes that this was exactly what happened in the US when 1930’s bank robbers used the new-fangled automobile to flee across state lines. The US solution was to make bank robbery (along with auto-theft and other related offences) into federal offences rather keeping them as state-specific infractions. However, this solution does not look to be practical for cyberspace, because there is no global body with the equivalent reach over the world’s countries that the US federal government had over the individual US states.

Others have argued for a specific law for cyberspace that is orthogonal to all national laws (the *Lex Mercatoria* from the beginning of the last millennium – an early attempt at a single market – is often cited as a historical example of such an approach²⁴). However, attempts at developing a *Lex Cyberspace* have, as with a super-federalist approach, foundered on the lack of institutions to sponsor it.

²⁴Sachs [117] argues that the documentary evidence from the period shows that merchants were substantially subject to local control and that the *Lex Mercatoria* did not actually exist as a uniform set of regulations for merchants, evolved by them and enforced by their own courts irrespective of the local jurisdiction. He says, ‘The traditional interpretation has been retained, not for its accuracy, but for ideological reasons and for its long and self-reinforcing pedigree’ and continues that he takes ‘no position on the merits of shielding multinational actors from domestic law’ but ‘merely denies that the Middle Ages provide a model for such policies.’

The practical approach that has been taken to deal with cross-jurisdictional criminals is to try and harmonise national laws within a consistent international framework. The relevant treaty for the specific harms (as listed above) that cannot be dealt with by existing ‘offline’ legislation is the 2001 Convention on Cybercrime [29] which sets out the required offences, provides the requisite definitions and sets out a uniform level of punishments.

All of the EU states have signed the convention, but some six years later only 12 (Bulgaria, Denmark, Estonia, France, Cyprus, Latvia, Lithuania, Hungary, the Netherlands, Romania, Slovenia and Finland) have ratified whereas 15 (Belgium, Czech Republic, Germany, Ireland, Greece, Spain, Italy, Luxembourg, Malta, Austria, Poland, Portugal, Slovakia, Sweden and the United Kingdom) have failed to ratify so far – usually because their law doesn’t yet cover particular issues (or tariffs are inadequate) rather than because of a complete lack of applicable law. If the harmonisation approach is to bear fruit, this process needs to be speeded up.

Recommendation 13: We recommend that the European Commission put immediate pressure on the 15 Member States that have yet to ratify the Cybercrime Convention.

The Convention has also been signed by a number of non-EU countries including Canada, Japan, South Africa, Ukraine and the United States. Of these only Ukraine and the United States have ratified. Quite clearly, the wider the adoption of the Convention the better.

In 2003 the Council of the European Union adopted a framework decision on attacks against information systems [30] which has subtly different definitions, and which distinguishes the offences of ‘illegal access to information systems’, ‘illegal system interference’, ‘illegal data interference’ along with ‘instigation, aiding, abetting and attempt’. These offences, along with some mandatory maximum (not minimum) tariffs, had to be in place by 2005.

In May 2007 the EU Commission issued a draft communication on cybercrime [47]. This defined cybercrime as traditional crimes committed over electronic networks, illegal content (child abuse pictures, etc), and ‘crimes unique to electronic networks’. The section on legislation was vague, suggesting legislation against ‘identity theft’ (which would surely already exist for offline theft), and ‘regulation on the responsibility of different actors in the relevant sector’ which is a content-free description. However, other public comments [98] suggested regulations could include mandatory blocking of sites containing bomb-making instructions and controls on search engines to prevent them returning results for words such as ‘bomb, kill, genocide or terrorism’.

8.2 Improving co-operation across jurisdictions

Co-operation across law enforcement jurisdictions is essential for online crime, yet there are very serious impediments against police forces working together.

8.2.1 Defining the problem

Given limited resources, police forces must make tough choices in deciding which crimes to investigate. In the case of electronic crime, one of the first questions raised is how many

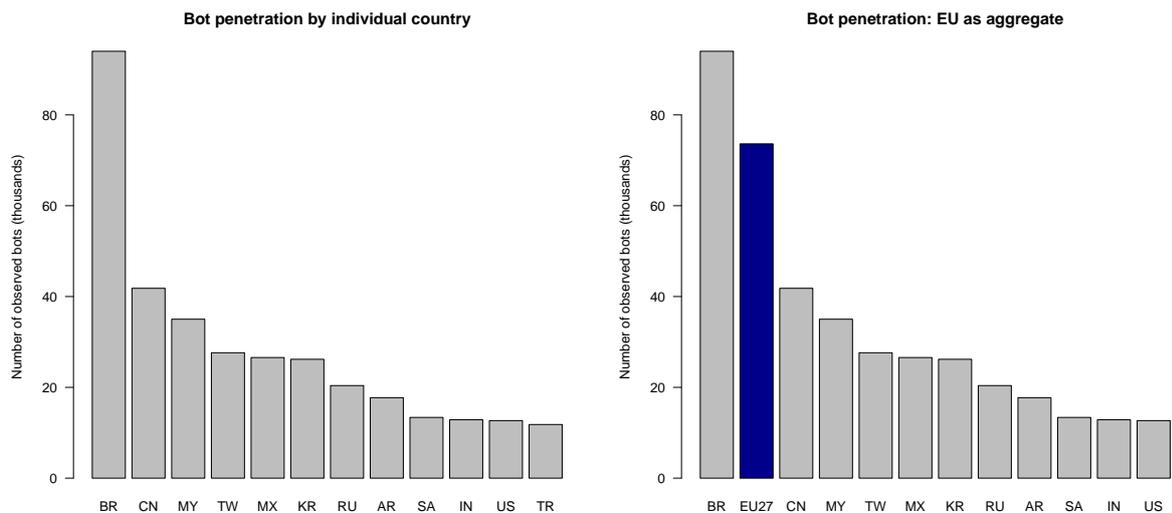


Figure 13: Not our problem? Number of global botnet victims identified by the Chinese HoneyNet Project between June 2006 and December 2007. No European country is among the top dozen alone (left), but the European Union as a whole is second only to Brazil (right). *Source:* Own aggregation based on data of [138]

of the country’s citizens are affected, and how many of the country’s computers are being used to launch attacks. Using this criteria, most attackers are not worth pursuing, even if (viewed as a whole) they are having a devastating effect (see Figure 13). Even in cases that are deemed worth pursuing, investigations invariably lead to computers located in other countries. International co-operation slows down investigations and drives up costs, even as it lessens the relevance to the country where the investigation began.

As a result, very few cyber-criminals are caught and successfully prosecuted. Lower risk levels in turn makes attacks more attractive and therefore more prevalent.

The fragmentation of law enforcement combined with the international nature of cyber-crime makes defender’s jobs harder as well. Banks have to allocate substantial resources to liaise with law enforcement agencies in many jurisdictions. These targets of cyber-crime then become less likely to pursue attacks involving distant or difficult jurisdictions.

8.2.2 Methods for co-operation

There are several traditional options for law enforcement agencies when they determine that a digital crime involves machines based in another country. Unfortunately, each is cumbersome and expensive.

Option 1: Increase funding for joint operations The first choice is to establish a *joint operation* between police forces. In a typical joint operation pursuing a cyber-crime, the country where the investigation began does most of the work while the co-operating country serves warrants and obtains evidence as requested by the originating country’s

force – this is a typical way of dealing with drug importation offences. A major difficulty with joint operations is that it is hard to predict what the cost will be prior to approving the co-operation. Joint operations are largely unfunded and carried out on a quid pro quo basis, so they cannot be relied upon as a fundamental response to all cyber-crimes. Nevertheless, increasing the funds available for supporting joint operations involving cyber crime is one policy option.

Option 2: Mutual legal assistance treaties Where joint operations are not possible, co-operation may still be possible via a mutual legal assistance treaty (MLAT). MLATs require a political decision taken by the requested country’s foreign ministry to determine whether co-operation can commence. While this is certainly feasible in most cases of cyber-crime (with the exceptions likely to be politically motivated crimes), MLATs are very slow to process. Hence, many investigators prefer to avoid using them where possible.

Essentially, the somewhat cumbersome requirements for international co-operation are largely acceptable for physical crimes, since cross-border activity is rare. In a digital environment where nearly all crimes cross borders, existing mechanisms do not suffice.

Option 3: Cyber-security co-operation using NATO as a model Quite clearly, more resources need to be devoted to tackling cross-border cyber crime. This requires cross-border co-operation with those who share the common cause – but cannot at present for reasons of sovereignty be done by cross-border policing actions.

The problem of countries working together for a common cause whilst preserving many aspects of their sovereignty has already been tackled by the military – whether it was SHAPE in World War II or NATO today. The model is that each country takes its own political decision as to what budget to set aside for fighting cyber crime. However, in all cases, one part of this budget is used to fund the presence of liaison officers at a central command centre. That command centre takes a European wide view of the problems that are to be tackled – and the liaison officers are responsible for relaying requests to their own countries and passing back the results as may be necessary.

This might be seen as a permanent ‘joint operation’ but it avoids the glacier-like speed of MLAT arrangements by insisting that liaison officers are able to immediately assess which requests carry no political baggage but can be expedited immediately.

Recommendation 14: We recommend the establishment of an EU-wide body charged with facilitating international co-operation on cyber crime, using NATO as a model.

9 Other issues

9.1 Cyber-insurance

Cyber-insurance has been cited by various authors as tool for cyber-risk management, in particular to transfer residual risk which cannot be mitigated with other types of security investment [74, 14, 92].

We define cyber-insurance as insurance contracts between insurance companies and enterprises or individuals covering financial losses incurred through damage or unavailability of tangible or intangible assets caused by computer or network-related incidents. This includes, inter alia,

- **first party risks:** destruction of property and data, network business interruption, cyber-extortion, cyber-terrorism, identity theft, recovery after virus or hacker attack;
- **third party risks:** network security liability, software liability, web content liability, intellectual property and privacy infringements due to data theft or loss.

One might expect the cyber-insurance market to be thriving, and a brisk market is generally acknowledged to be socially beneficial for four reasons.

1. **Incentives to implement good security.** Insurance companies may differentiate premiums by risk classes so that insured parties who take appropriate precautions will pay lower premiums. In theory, this should reward effective safeguards and go some way to mitigating the agency effects that often lead to security measures being deployed for mere due-diligence and directors' peace of mind. Insurers will also assign different software products and management practices to different risk classes, thus passing on pressure to develop secure products to the software industry (assuming that markets are competitive).

However, practice looks a bit different. While banks buying nine-figure cover were actually inspected, firms purchasing more modest policies typically find their premiums based on non-technical criteria such as firm size or individual loss history. Some exceptions include Chubb, who offers rebates to firms that test their security systems regularly [25]. Also the differentiation between off-the-shelf and customised software is common (standard software is considered more secure and thus rewarded with lower premiums). We are not aware of any differentiation between operating systems, probably because there is little variation in the clients' installed base.

2. **Incentives for security R&D.** As part of their risk management, insurers gather information about the risks they are underwriting, and the claims history is particularly relevant. The more business they underwrite, the better they are informed, the more accurately premiums can be calculated and the more competitive they become. To bootstrap this virtuous circle, insurers have an incentive to reinvest part of their revenues to improve their knowledge base. European insurers say that they are investing in research, both via in-house engineers and in co-operation with security technology firms. (We are aware though of only one concrete case in which an insurance association funded original research on the vulnerabilities in a system.)

3. **Smooth financial outcome.** As for all insurance contracts, insured parties exchange uncertainty about their future assets for a fixed present cost. This reduces the variance of their asset value over time. They can re-allocate capital to their core business instead of holding extra reserves to self-insure against IT failures. This is particularly useful for industries that depend on IT, but do not consider it as their core activity.
4. **Market-based security metric.** As discussed earlier in Section 4.2.5 of this report, insurance premiums may serve as market-driven metrics to quantify security. This metric fits well in an investment-decision framework, as risk managers can weigh the costs of security investment against reductions in insurance premiums [74]. Indeed, the insurers' actual claims history would be an extremely valuable source of data for security economists, but insurers consider this to be highly sensitive because of the competitive advantage derived from better loss information.

That at least is the theory; it makes cyber-insurance sound compelling. Yet the market appears to perform below expectations. The USD 350 million estimated global market size in 2005 [31] is only one-tenth of a forecast made for 2005/06 by the Insurance Information Institute in 2002 [88] and below one fifth of a revised forecast from 2003 [82]. According to the 2007 CSI Computer Crime and Security Survey, only 29% of the large US-based companies surveyed reported having any insurance policy covering cyber-risks. This is around the same share as in previous years²⁵ and in line with the judgement of industry experts in Europe.

In fact, the cyber-insurance market has long been somewhat of an oddity. Until Y2K, most companies got coverage for computer risks through their general insurance policy, which typically covered losses due to dishonesty by staff as well as theft by outsiders. There were also some specialist markets, particularly for banks who wanted substantial coverage. A typical money-center bank in the late 20th century carried USD 200 million of 'Bankers Bond and Computer Crime' cover, in which market Lloyds of London was the dominant player. Banks purchasing these policies had to have their systems assessed by specialist information security auditors and coverage was typically conditional on the remediation of known problems. Premiums were typically 0.5% of the sum assured in the 1980s, and about 1% in the 1990s (following a claim). In the run-up to Y2K, many UK and US insurers stopped covering computer risks; the market resumed in 2002–2004 with premiums initially well above 1%. Competition has pushed these down to the range of 0.3–0.5%.

In the German market, TELA, an insurance subsidiary of Siemens, started underwriting IT risks (including software risks) in the 1970s. It was sold to Allianz in 2001 and, in the aftermath of 9/11, Allianz discontinued TELA's cyber-insurance product line. Y2K has been exempted from coverage, but there is no sign that insurers stopped covering computer risks in general. Allianz returned to the cyber-insurance market in 2004 (dropping the name TELA) but found that subsidiaries of its international competitors filled the gap in the German cyber-insurance market. TELA had a loss research department until 1988, before it was hived off in 1988 as Tescon, which became an independent security consultancy in 2002.

²⁵28% in 2005, 25% in 2006 [28]

Some industry sources blame a lack of good actuarial data for the slow adoption rate, but this would not explain the flat trend over several years. An alternative explanation is that losses from some information security risks are highly correlated globally, which makes cyber-insurance uneconomical. There are basically two types of risk: risks local to an insured company, for example that a financial manager commits a large fraud by abusing his computer access, or that a specific vulnerability is exploited by an outsider as in the Levin case; and global risks, for example that the firm loses several days' trading because of an attack by a virus or worm. Homogeneity in installed platforms means that attacks of the second kind can spread to millions of systems within minutes. This points to a link between diversity and insurability: correlated risks require additional safety premiums that render cyber-insurance policies too expensive for large parts of the market [14]. Other demand-side barriers to cyber-insurance include a lack of awareness among insurance brokers, risk managers and senior executives; the uncertainty about accountability for cyber-crime losses; the difficulty of pricing such losses; and the absence in some industries of industry standards [31].

German industry experts whom we interviewed when preparing this report were most wary of cumulated risks. In fact, they claimed to find little evidence of correlation in their (more or less long) historical data. This could be due to a lack of statistical power or a result of specific exclusions designed to keep correlated risks out of the portfolio. Typical steps to avoid correlation include excluding damage incurred by untargeted attacks or limiting coverage when the insureds' suppliers dump liability (e.g. by waving right of recourse agreements). Clients dislike these exclusions and even occasionally name them as reasons for deciding not to buy cyber-insurance. This shows the interdependence of diversity (Section 7), liability (Section 6) and cyber-insurance. But cyber-insurance is also related to information asymmetries (Section 4).

Industry experts reckon that the lack of awareness of cyber-risks is the most important demand-side barrier, whereas they consider the elasticity of demand to premium changes very low. However, they observe that some European clients have started to take notice of media reports of US breaches. A comprehensive breach-disclosure law for the EU might help overcome the slack in demand for cyber-insurance.

Government action to overcome these barriers and help establish a wider market for cyber-insurance could be justified against the backdrop of expected gains in social welfare due to positive externalities arising from a viable market for cyber-insurance. Several options are conceivable in theory.

Option 1: Compulsory cyber-insurance One option could be to make insurance compulsory for networked PCs, just as every car that runs on Europe's roads must be insured. This would certainly spur demand for cyber-insurance, but policy makers must be very careful here. The insurance market for firms appears to have few claims and high premiums, and whether this is ascribed to risk correlation or simply lack of competition, making such products a compulsory purchase would be seen as an unjustified tax and furthermore one that lined the pockets of an industry that contributes little directly to the solution of cyber-security problems. The opponents of such a tax would see this tax as a deadweight on competitiveness and productivity growth; and they would point to the current lack of claims against owners of infected machines or their ISPs. Although our Recommendation 4, of a fixed penalty charge, will if adopted cause claims to appear,

it would probably be best to wait and see how the market copes with that.

A transition from a world without much cyber-insurance to a regime with full cyber-insurance coverage is of course possible in the longer term, and this may happen by sectors. The criticality of an application is a good criterion for selecting sectors for compulsory insurance; and a particularly strong case can be made where actors of limited means have the ability to cause substantial damage (this is the essence of the case for mandatory car insurance). Taking transition dynamics into account, another criterion could be the growth rate of IT dependence (as it is more difficult to replace existing systems with insurable ones than building an insurable infrastructure from scratch). Compulsory cyber-insurance might also be targeted at those market segments that are least likely to thrive under their own steam, such as volume contracts for small and unlikely losses [14], or against events for which large enterprises would prefer self-insurance over risk transfer [16], though then the regulator might be accused of unduly favouring the insurance industry.

Option 2: Government re-insurance Secondary coverage for conventional insurance business is supplied by just a few re-insurers, which try to balance undue concentration of risk through global diversification. However, globally-connected networks and cross-border crime mean that cyber-risks are hard to hedge geographically. Primary insurance companies started to explicitly exclude cyber-risks from existing contracts in January 2002, because their reinsurance companies were concerned about a global ‘cyber-hurricane’, which they would not be able to deal with [35]. The market cycle has now turned and re-insurance for cyber-risks is available on reasonable conditions. But this may change over time, in particular if the volume grows as the market matures and re-insurance is sought for larger chunks of (possibly correlated) cyber-risk. If this turns out to become a constraining factor, governments might be asked to step in.

While government re-insurance can create insurance markets where otherwise there would be no supply, such measures must be carefully designed to avoid a regime in which profits are private (to the insurers’ shareholders), losses are socialised (born by the taxpayer), and systems remain insecure (because the government intervention removes the incentive to build properly secure products). Again, one must bear in mind the potential for government reinsurance to be seen as undue state aid. There are circumstances in which it might be sought as a temporary measure to steady the market or specific sectors of it. But it must be set up with sunset provisions so that it can be gradually reduced and replaced by private coverage. In the meantime, if information sharing is properly dealt with by the regulation, the state could have access to detailed claims data and would have the opportunity to understand the real effects of cyber-risks on businesses in much more detail than at present.

Option 3: Additional anti-discriminatory regulation Policy makers might be tempted to support fair access to insurance products by requiring insurers to cap premiums or charge fixed premiums. The political pressure to do so would likely rise if the insurance product were compulsory or partly backed with state re-insurance. For example, the public-private partnership of natural catastrophe insurance in France [93] includes provisions for state-regulated premiums. However, premium differentiation is the key to creating incentives for good security. If bad security practices are not penalised by higher premiums, people may even act more riskily – as with some government-backed flood in-

insurance programs, which fostered construction on flood-prone river banks by guaranteeing insurance coverage at fixed premiums.

Option 4: Financial instruments for risk sharing Correlated risks might be dealt with by risk transfer to, and diversification on, broader financial markets. Specially designed financial instruments could allow insurers to pass on packages of well-defined risk to other market participants in exchange for a risk premium. *Exploit derivatives* (see Section 4.2.5) are vehicles for insurers to hedge against the discovery of vulnerabilities that cause significant loss events across their portfolios. The insurers' cyber-risk managers would develop models to map the expected actual loss amounts to a portfolio of exploit derivatives taking into account their clients' risk profiles in terms of software installed and assets at risk. *Cat bonds* [33], another class of instruments for insurance risk securitization, do not require this mapping. Their pay-out function is defined on actual impact rather than on the theoretical possibility of a breach. Both types of instruments allow dealing in cumulated risk – at least to a certain extent – because market participants can diversify their investment between asset classes.

There is some experience with cat bonds in flood and natural-disaster insurance, but no experience at all with exploit derivatives, as the latter are more specific to IT. A difficulty in applying cat bonds to IT might lie in the moral hazard problem: speculators might find themselves in situations where causing or commissioning a cyber-attack would improve their financial wealth. Conventional insurance can deal with moral hazard by strictly limiting cat bond pay-out functions to purely natural perils.

Option 5: Insurable infrastructure design The interdependent nature of cyber-risk means that insurability and incentives to buy insurance are determined by the technical environment, such as network topology, configuration and protocols [90, 108, 24, 14, 16, 17]. While Bolot and Lelarge's recommendation:

‘[N]etwork algorithms and network architecture might be designed or re-evaluated according to their ability to help implement desirable economic policies, such as the deployment of insurance’ [17]

remains rather vague, concrete measures to improve insurability can be taken by increasing diversity. For example, an ISP that was totally dependent on Cisco routers should logically pay higher premiums than one which had diversified by purchasing Juniper equipment as well. Formal economic models show that equilibrium premiums for diverse systems are below those of homogeneous ones even if the unconditional probability of failure of each diverse node is higher than the unconditional probability of failure of the homogeneous nodes [14]. System diversity should be a policy maker's goal not only for reasons of fair competition but also to increase robustness and resilience.

Conclusions on cyber-insurance If we order the options by priority, then the ideal long-term goal is building an insurable infrastructure, or at least seeing to it that insurability is not harmed by infrastructure design. Second, better financial instruments to facilitate risk transfer would be useful; policy makers should ensure that their use isn't impeded by the regulatory and supervision framework. Making cyber-insurance compulsory would be a heavyweight intervention in an immature market and should therefore

be avoided. (It might though be a workable last resort in specific sectors should they come under pressure from cyber-incidents in the future.) The provision of government re-insurance is expensive and is rather likely to create misaligned incentives; and premium differentiation is so essential for cyber-insurance that any attempt to fix or influence premiums should be strictly avoided.

We conclude the chapter on cyber-insurance without a straight recommendation because we see that the market is becoming more and more competitive over time. And we believe that some of our other recommendations, if properly implemented, will help the market to develop anyway. In particular, breach-disclosure legislation will raise awareness and thus help to overcome the most important demand-side barrier. Better statistics should help insurers to improve their actuarial models, more diversity might reduce risk correlation, and fixing liability may help rid insurance contracts of the more vexatious exclusions. (Despite this tentatively optimistic outlook, the European dimension of the cyber-insurance market is an area where more policy-related research is needed, as most empirical data and literature focuses on the US market only.)

9.2 Security research and legislation

Security research is important, and occurs at a number of places in the value chain. First, blue-sky (typically academic) researchers think up new algorithms, protocols, operating-system access-control schemes and the like. Second, applied researchers investigate how particular types of systems fail, and devise specific proposals for submission to standards bodies. These researchers can be academic, industrial, or a mix. Third, research and development engineers produce prototypes and write code for specific products and services. Fourth, users of these products or services discover vulnerabilities. These are often design or implementation errors rather than flaws in the underlying security technology. Examples of design issues include protocol failures, while implementation errors consist largely of programming mistakes such as buffer overflows and race conditions.

Public policy has got in the way of security research on a number of occasions. The debate on cryptography policy during the 1990s led to EC Regulation 1334/2000 on Dual Use Goods under which the export of cryptographic software in intangible form (e.g. researchers swapping source code) became subject to export control. There are a number of exemptions and open licenses that researchers can use; for example, material in the public domain is generally exempt, and many countries have open general licenses for the export of standard cryptographic software. However, open licenses may come with registration requirements or limitations on the length of cryptographic keys. There is also a Community General Export Authorisation (CGEA), but to use this firms in some countries have to register with its national export control body.

One problem is that many small software developers are unaware of this control regime and may be technically in breach of its implementation provisions in some Member States. The export-licensing regime in a country such as the UK is aimed at large companies that export armaments, and is simply not organised to communicate with tens of thousands of small specialist companies whose business may be electricity meters or engine controllers, and who incorporate cryptography as a subsystem. Export controls are of concern to academics (though to a small extent, as most crypto researchers can simply place their work product in the public domain); of some concern to applied researchers; and of most

concern to research and development staff at companies.

Unfortunately, the research conducted by the Commission into the working of the Regulation appears to have spoken only to the large firms who are already aware of the controls, and to their trade associations [83]. Even so it discovered a number of non-tariff barriers, for example in inconsistent national implementations that resulted in goods being rerouted circuitously to avoid controls. We recommend that ENISA get engaged in the process of reforming the export control regime. There is a specific problem relating to cryptography in the 56-bit keylength restriction for general licenses. This was introduced when the encryption standard was DES with 56-bit keys; but since DES can be broken in hours the world standard is now AES with 128-bit or 256-bit keys. Hence many more security products will in theory require export licensing (as will other products containing some limited cryptographic functionality). It is not in the interests of the information security community to go unrepresented as these regulations evolve.

A more recent concern is that in some Member States, well-meant but poorly drafted legislation has impeded security research. Although Conventions, Directives and Decisions use language such as ‘without right’, the national transposition or implementation has often interpreted this poorly. In Germany, the criminal law code (Strafgesetzbuch) has been amended with a new section 202c that makes it an offence to produce, supply, sell, transmit, publish or otherwise make accessible any password, access code or software designed to perpetrate a computer crime, in preparation for such a crime. This has been opposed as excessive by many researchers who see it as threatening those who possess system engineering tools for innocuous purposes [4]. In the UK, the Government amended the Computer Misuse Act to make it an offence to ‘supply or offer to supply, believing that it is likely to be used to commit, or to assist in the commission of [a computer offence]’ so that it is the meaning of ‘likely’ which will determine whether an offence has been committed. The government’s response to concern about the circumstances in which an offence would be committed has been to promise to publish guidance for prosecutors as to when the law should or should not be invoked.

In both cases the concern is that IT and security professionals who make network monitoring tools publicly available or disclose details of unpatched vulnerabilities could be prosecuted. Indeed, most of the tools on a professional’s laptop, from `nmap` through `wireshark` to `perl` could be used for both good and bad purposes. The resulting legal uncertainty has a chilling effect on security research [26].

Recommendation 15: We recommend that ENISA champion the interests of the information security sector within the Commission to ensure that regulations introduced for other purposes do not inadvertently harm security researchers and firms.

Although the two most harmful regulations up till now have been in the areas of export control and cyber-crime, there will no doubt be more. The industry needs an advocate in Brussels to ensure that its interests are taken into account when directives and regulations are being formulated – and as they evolve over time. In the case of export control, we recommend that ENISA push for cryptography to be removed from the dual-use list. In the case of dual-use tools that can be used for hacking as well as for bona-fide research and administrative tasks, we recommend ENISA take the position that sanctions should only apply in the case of ill intent.

10 Conclusions

As Europe moves online, information security is becoming increasingly more important: first, because the direct and indirect losses are now economically significant; and second, because growing public concerns about information security hinder the development of both markets and public services.

While information security touches on many subjects from mathematics through law to psychology, some of the most useful tools for both the policy analyst and the systems engineer come from economics. Systems often fail not for some technical reason, but because the incentives were wrong. (Indeed, incentive failures often underlie technical failures.) As a result, security economics has become a live subject of research over the last seven years.

In this report, we have provided an analysis based on security economics of the practical problems in network and information security that the European Union faces at this time. We have come up with fifteen policy proposals that should make a good next step in tackling the problems. We therefore hope that they will provide the basis for constructive action by ENISA and the European Commission in the future.

Acknowledgements

The authors are grateful to acknowledge input from the attendees at the ENISA stakeholders' meeting on December 10th, 2007; from people in the security industry who talked to us, mostly off the record; from colleagues in the security groups at Cambridge and Dresden; from members of the Advisory Council of the Foundation for Information Policy Research, particularly Nick Bohm, Alan Cox, Douwe Korff, Jim Norton and Martyn Thomas; and from Alexander Korff of Clifford Chance. Responsibility for any errors and omissions of course remains ours alone.

Ross Anderson, *Cambridge University*
Rainer Böhme, *Technische Universität Dresden*
Richard Clayton, *Cambridge University*
Tyler Moore, *Cambridge University*

January 29th 2008

List of Figures

1	Web trojan generator interface and data theft crimeware interface	10
2	Crimeware affiliate marketing	11
3	US breach disclosure statistics	23
4	Distribution of breach reports across sectors and breach types	24
5	US breach disclosures by sector and type	25
6	Severity by breach type and sector	25
7	Phishing-site lifetimes per bank	28
8	Average annual reported losses per enterprise attributed to computer crime	29
9	Proportion of annual reported losses attributed to different threat categories	30
10	Denial-of-service attacks per week in 2004 and 2005	33
11	Global attack tactics: Phishing sites vs malware	36
12	Patch-development times for different operating systems	61
13	Number of global botnet victims identified by the Chinese HoneyNet Project	80

List of Tables

1	Proportion of top threats that undermine confidentiality according to different Symantec reports	32
2	Recommended breakdowns for information security indicators	35
3	Barriers to sharing security information	41
4	Time after public disclosure of vulnerabilities before a patch is issued and an exploit is published	62
6	Information society indicators: security problems of individuals	104
7	Information society indicators: security problems of enterprises	105
8	Information society indicators: security problems of large enterprises	106
9	Information society indicators: individuals' use of security technology . . .	107
10	Information society indicators: enterprises' use of security technology . . .	108
11	Information society indicators: enterprises' update practices	109
12	Information society indicators: perceived barriers to e-commerce	110
13	Internet exchange point sizes	111

List of Acronyms

AES	Advanced Encryption Standard
APACS	Association of Payment and Clearing Services (UK)
APWG	Anti Phishing Working Group
ATM	Automatic Teller Machine
B2B	Business-to-business
B2C	Business-to-consumer
BSI	Bundesamt für Sicherheit in der Informationstechnik (DE)
CERT/CC	Computer Emergency Response Team (Co-ordination Center)
CGEA	Community General Export Authorisation
CIWIN	Critical Infrastructure Warning Information Network
CNI	Critical National Infrastructure
DDoS	Distributed Denial of Service
DES	Data Encryption Standard
DG	Directorate General
DRM	Digital Rights Management
EAL	Evaluation Assurance Level (of the Common Criteria standard)
EMV	Europay, Mastercard and VISA (a standard for chip card payment systems)
ENISA	European Network and Information Security Agency
EPCIP	European Programme for Critical National Infrastructure Protection
EULA	End-user License Agreement
GCHQ	Government Communications Headquarters (UK)
GPS	Global Positioning System
ICT	Information and Communication Technology
IDABC	Interoperable Delivery of European eGovernment Services to Public Administrations, Businesses and Citizens
IDS	Intrusion Detection System
IP	Internet Protocol
IRC	Internet Relay Chat
ISAC	Information Sharing and Analysis Centre
ISP	Internet Service Provider
ISTR	Internet Security Threat Report
IT	Information Technology
IXP	Internet Exchange Point
LHS	Left-hand scale
MLAT	Mutual Legal Assistance Treaty
MLS	Multilevel Secure
MSSP	Managed Security Service Provider
NACE	Nomenclature of Economic Activities (an industry classification)
NATO	North Atlantic Treaty Organisation
NGO	Non-governmental Organisation
NIS	Network and Information Security
NSP	Network Service Provider
ODF	Open Document Format
OECD	Organisation for Economic Co-operation and Development

OEM	Original Equipment Manufacturer
OMA	Open Mobile Alliance
OS	Operating System
PDA	Personal Digital Assistant
PIN	Personal Identification Number
PISCE	Partnership for ICT Security Incident and Consumer Confidence Information Exchange
PKI	Public Key Infrastructure
PNG	Portable Network Graphics
R&D	Research and Development
RFID	Radio Frequency Identification
RHS	Right-hand scale
SCADA	Supervisory Control and Data Acquisition
SIM	Subscriber Identity Module
SME	Small and Medium Enterprise
WGA	Windows Genuine Advantage (an anti-piracy tool)

References

- [1] Abu El-Ata, A., Aeberhard, M., Furrer, F. J., Gardiner-Smith, I. and Kohn, D. (2002): *Our PKI Experience, Towards and Enterprise-wide PKI: Concepts, Architecture, and Decision Drivers for the CSG-PKI*. Syslogic Press, Switzerland
- [2] Acquisti, A., Friedman, A. and Telang, R. (2006): Is there a cost to privacy breaches? An event study. *Workshop on the Economics of Information Security (WEIS)*, Univ. of Cambridge, UK. <http://weis2006.econinfosec.org/docs/40.pdf> (last access: 13 Nov 2007)
- [3] Akerlof, G. A. (1970): The market for ‘lemons’: quality uncertainty and the market mechanism. *Quart. J. Economics*, **84**, 488–500
- [4] Anderson, N. (2007): German ‘anti-hacker’ law forces hacker sites to relocate. *Ars Technica* (14 August), at <http://arstechnica.com/news.ars/post/20070814-german-anti-hacker-law-forcing-hacker-sites-to-relocate.html>
- [5] Anderson, R. J. (2001): *Security Engineering – A guide to building dependable distributed systems*, Wiley. <http://www.cl.cam.ac.uk/~rja14/book.html>
- [6] Anderson, R. J., Moore, T. W. (2007): Information security economics – and beyond. *Advances in Cryptology – Crypto 2007*, LNCS 4622, Springer Verlag, Berlin Heidelberg, 68–91 http://www.cl.cam.ac.uk/~rja14/Papers/econ_crypto.pdf (last access: 10 Jan 2008)
- [7] Anti-Phishing Working Group: <http://www.antiphishing.org/>
- [8] APACS (2007): Card fraud losses continue to fall. Press Release, 14 March. http://www.apacs.org.uk/media_centre/press/07_14_03.html
- [9] Arora, A., Krishnan, R., Telang, R., Yang, Y. (2005): An empirical analysis of vendor response to disclosure policy. *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA.
- [10] Bangeman, E. (2007): The insanity of France’s anti-file-sharing plan: L’État, c’est IFPI, *Ars Technica* (25 November), at <http://arstechnica.com/news.ars/post/20071125-the-insanity-and-genius-of-frances-anti-file-sharing-plan.html>
- [11] BBC (2007): Devices attached to cash machines, (15 October), at <http://news.bbc.co.uk/1/hi/england/cambridgeshire/7044894.stm> (last access: 10 Jan 2008)
- [12] Beattie, S., Arnold, S., Cowan, C., Wagle, P., Wright, C., Shostack, A. (2002): Timing the application of security patches for optimal uptime, *Proc. of LISA 2002*, 233–242

- [13] Bohm, N., Brown, I. and Gladman, B. (2000): Electronic commerce: Who carries the risk of fraud? *J. Information, Law and Technology*, **3**, at http://www2.warwick.ac.uk/fac/soc/law/elj/jilt/2000_3/bohm/
- [14] Böhme, R. (2005): Cyber-insurance revisited. *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA. <http://infoecon.net/workshop/pdf/15.pdf> (last access: 13 Nov 2007)
- [15] Böhme, R. (2006): A comparison of market approaches to software vulnerability disclosure. In Müller (ed.): *Emerging Trends in Information and Communication Security (ETRICS)*, LNCS 3995, Springer Verlag, Berlin Heidelberg, 298–311
- [16] Böhme, R., Kataria, G. (2006): Models and measures for correlation in cyber-insurance. *Workshop in the Economics of Information Security (WEIS)*, Univ. of Cambridge, UK. <http://weis2006.econinfosec.org/docs/16.pdf> (last access: 30 Nov 2007)
- [17] Bolot, J. and Lelarge, M. (2007): A new perspective on Internet security using insurance. *INRIA Research Report*. <http://hal.inria.fr/docs/00/18/14/39/PDF/cyber-RR.pdf> (last access: 13 Nov 2007)
- [18] Brenner, S (2007): Bonnie & Clyde and cybercrime. <http://cyb3rcrim3.blogspot.com/2007/11/bonnie-clyde-and-cybercrime.html> (last access: 10 Jan 2008)
- [19] California State Senate (2002): Assembly Bill 700. http://info.sen.ca.gov/pub/01-02/bill/asm/ab_0651-0700/ab_700_bill_20020929_chaptered.pdf
- [20] Campbell, K., Gordon, L. A., Loeb, M. P. and Zhou, L. (2003): The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *J. Computer Security*, **11** (3), 431–448
- [21] Casper, C. (2007): Examining the feasibility of a data collection framework. *ENISA Technical Report*.
- [22] Cavusoğlu, H., Cavusoğlu, H., Zhang, J. (2006): Economics of patch management. *Workshop in the Economics of Information Security (WEIS)*, Univ. of Cambridge, UK. <http://weis2006.econinfosec.org/docs/5.pdf> (last access: 10 Jan 2008)
- [23] Cavusoğlu, H., Mishra B. and Rangunathan, S. (2004): The effect of Internet security breach announcements on market value: Capital market reactions for breached firms and Internet security developers. *Int. J. Electronic Commerce*, **9** (1), 69–104
- [24] Chen, P.-Y., Kataria, G. and Krishnan, R. (2005): Software diversity for information security. *Workshop in the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA. <http://www.infoecon.net/workshop/pdf/47.pdf> (last access: 3 Dec 2007)
- [25] Chubb (2007): Chubb encourages adoption of new information security best practices. Press Release, 3 December. <http://www.chubb.com/corporate/chubb7880.html>

- [26] Clayton, R. (2007): Hacking tools are legal for a little longer. <http://www.lightbluetouchpaper.org/2007/06/19/hacking-tools-are-legal-for-a-little-longer/>
- [27] Collins, B. St. J. (1995): Unfair terms in consumer contracts regulations 1994, at <http://webjcli.ncl.ac.uk/articles3/collins3.html>
- [28] Computer Security Institute (2007): *The 12th Annual Computer Crime and Security Survey*. <http://www.gocsi.com/> (last access: 8 Nov 2007)
- [29] Council of Europe (2001): Convention on Cybercrime, CETS 185. <http://conventions.coe.int/Treaty/Commun/QueVoulezVous.asp?NT=185&CL=ENG> (last access: 10 Jan 2008)
- [30] Council of the European Union (2003): Council Framework Decision on attacks against information systems. <http://register.consilium.eu.int/pdf/en/03/st08/st08687-re01en03.pdf> (last access: 10 Jan 2008)
- [31] Critical Infrastructure Protection Program (2007): Cyber insurance. *The CIP Report*, 6 (3). http://cipp.gmu.edu/archive/cip_report_6.3.pdf (last access: 3 Dec 2007)
- [32] Dacey, R. (2003): Information security: Progress made, but challenges remain to protect federal systems and the nation's critical infrastructures. US General Accounting Office (GAO), GAO-03-564T (April) 1-75
- [33] D'Arcy, S. and France, V. G. (1992): Catastrophe futures: A better hedge for insurers. *J. Risk and Insurance*, 59 (4), 575-600
- [34] Dawes, S.S., Birkland, T., Tayi, G.K. and Schneider, C.A. (2004): *Information, Technology, and Coordination: Lessons from the World Trade Center Response*. Center for Technology in Government. http://www.ctg.albany.edu/publications/reports/wtc_lessons (last access: 10 Jan 2008)
- [35] Duffy, D. (2002): Safety at a premium. *CSO Magazine*, December. <http://www.csoonline.com/read/120902/safety.html> (last access: 3 Dec 2007)
- [36] ECMA International (2006): ECMA-376 Office Open XML File Formats (OOXML) Standard.
- [37] Edelman, B. (2004): 180solutions installation methods and license agreement. <http://www.benedelman.org/spyware/180-affiliates/installation.html> (last access: 18 Dec 2007)
- [38] Edelman, B. (2007): Advertisers using WhenU. <http://www.benedelman.org/spyware/whenu-advertisers/> (last access: 18 Dec 2007)
- [39] Edelman, B. (2007): Spyware: Research, testing, legislation, and suits. <http://www.benedelman.org/spyware/> (last access: 18 Dec 2007)

- [40] Edelman, B. (2007): Zango practices violating Zango's recent settlement with the FTC. <http://www.benedelman.org/spyware/zango-violations/> (last access: 18 Dec 2007)
- [41] van Eeten, M. J. G. et al. (2007): *The Economics of Malware: Security Decisions, Incentives and Externalities*. Draft OECD report.
- [42] Edquist, C. and Hommen, L. (2000): Public technology procurement and innovation theory. In Edquist et al. (ed.): *Public Technology Procurement and Innovation*, Kluwer, Boston, MA
- [43] Emigh., A. (2006): The crimeware landscape: Malware, phishing, identity theft and beyond. *Anti-Phishing Working Group Technical Report*. http://www.antiphishing.org/reports/APWG_CrimewareReport.pdf (last access: 20 December 2007)
- [44] European Commission (2005): Green Paper on a European Programme for Critical Infrastructure Protection. COM(2005) 576 final. http://eur-lex.europa.eu/LexUriServ/site/en/com/2005/com2005_0576en01.pdf (last access: 10 Jan 2008)
- [45] European Commission (2006): i2010 Benchmarking Framework http://ec.europa.eu/information_society/eeurope/i2010/docs/benchmarking/060220_i2010_Benchmarking_Framework_final_nov_2006.doc (last access: 8 Nov 2007)
- [46] European Commission (2006): Report on the outcome of the Review of the EU regulatory framework for electronic communications networks and services in accordance with Directive 2002/21/EC and summary of the 2007 reform proposals. http://ec.europa.eu/information_society/policy/ecom/doc/library/proposals/com_review_en.pdf (last access: 22 Nov 2007)
- [47] European Commission (2007): Defining the Commission's global policy on the fight against cyber crime. Press Release IP/07/689. <http://www.europa.eu/rapid/pressReleasesAction.do?reference=IP/07/689>
- [48] European Communities (2006): French administration opts for OpenOffice. 7 July 2006. <http://ec.europa.eu/idabc/en/document/5695/469>.
- [49] European Economic Community (1985): Council Directive of 25 July 1985 on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products (85/374/EEC)
- [50] European Union (1993): Directive 93/13/EEC of 5 April 1993 on Unfair Terms in Consumer Contracts. http://eur-lex.europa.eu/smartapi/cgi/sga_doc?smartapi!celexapi!prod!CELEXnumdoc&lg=EN&numdoc=31993L0013&model=guichett
- [51] European Union (1997): Directive 97/7/EC of the European Parliament and of the Council of 20 May 1997 on the protection of consumers in respect of distance contracts. <http://europa.eu.int/eur-lex/lex/LexUriServ/LexUriServ.do?uri=CELEX:31997L0007:en:HTML>

- [52] European Union (1999): Directive 1999/93/EC of the European Parliament and of the Council of 13 December 1999 on a Community framework for electronic signatures <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31999L0093:EN:HTML>
- [53] European Union (2000): Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on certain legal aspects of information society services, in particular electronic commerce, in the Internal Market (Directive on electronic commerce) <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32000L0031:EN:HTML>
- [54] European Union (2002): Directive 2002/19/EC of the European Parliament and of the Council of 7 March 2002 on access to, and interconnection of, electronic communications networks and associated facilities (Access Directive) <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0019:EN:HTML>
- [55] European Union (2002): Directive 2002/20/EC of the European Parliament and of the Council of 7 March 2002 on the authorisation of electronic communications networks and services (Authorisation Directive) <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0020:EN:HTML>
- [56] European Union (2002): Directive 2002/21/EC of the European Parliament and of the Council of 7 March 2002 on a common regulatory framework for electronic communications networks and services (Framework Directive) <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0021:EN:HTML>
- [57] European Union (2002): Directive 2002/22/EC of the European Parliament and of the Council of 7 March 2002 on universal service and users' rights relating to electronic communications networks and services (Universal Service Directive) <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0022:EN:HTML>
- [58] European Union (2002): Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0058:EN:HTML>
- [59] European Union (2004): Public procurement: Commission examines discriminatory specifications in supply contracts for computers in four Member States, 13 October. <http://europa.eu/rapid/pressReleasesAction.do?reference=IP/04/1210&format=HTML&aged=0&language=EN&guiLanguage=en>
- [60] European Union (2006): Directive 2006/24/EC of the European Parliament and of the Council of 15 March 2006 on the retention of data generated or processed in connection with the provision of publicly available electronic communications services or of public communications networks and amending Directive 2002/58/EC <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2006:105:0054:01:EN:HTML>

- [61] European Union (2006): Directive 2006/123/EC of the European Parliament and of the Council of 12 December 2006 on services in the internal market <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2006:376:0036:0068:EN:PDF>
- [62] European Union (2007): Directive 2007/64/EC of the European Parliament and of the Council of 13 November 2007 on payment services in the internal market amending Directives 97/7/EC, 2002/65/EC, 2005/60/EC and 2006/48/EC and repealing Directive 97/5/EC Text with EEA relevance, <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2007:319:0001:01:EN:HTML>
- [63] Ettredge, M. and Richardson, V. J. (2002): Assessing the risk in e-commerce. *Proc. of the 35th Hawaii International Conference on System Sciences*, IEEE Press, Los Alamitos, CA
- [64] Fama, E. (1970): Efficient capital markets: A review of theory and empirical work. *J. Finance*, **25** (2), 383–417
- [65] Federal Trade Commission (2006): ChoicePoint settles data security breach charges; to pay \$10 million in civil penalties, \$5 million for consumer redress. Press Release. <http://www.ftc.gov/opa/2006/01/choicepoint.shtm> (last access: 22 Nov 2007)
- [66] Fox News (2008): Estonia Charges Solo Hacker for Crippling Cyberattacks. Jan 25 2008. <http://www.foxnews.com/story/0,2933,325547,00.html>
- [67] Franklin, J., Perrig, A., Paxon, V. and Savage, S. (2007): An inquiry into the nature and causes of the wealth of Internet miscreants. *Proc. of ACM CCS*, 375–388
- [68] Fraunhofer Institute for Systems and Innovation Research (2005): Innovation and public procurement. Review of issues at stake. Study commissioned by the European Commission. ftp://ftp.cordis.lu/pub/innovation-policy/studies/full_study.pdf (last access: 9 Nov 2007)
- [69] Galetsas, A. (2007): *Statistical Data on Network Security*. European Commission, DG Information Society and Media. ftp://ftp.cordis.europa.eu/pub/ist/docs/trust-security/statistics-network-security-050307_en.pdf (last access: 5 Nov 2007)
- [70] Gal-Or, E. and Ghose, A. (2005): The economic incentives for sharing security information. *Information Systems Research*, **16** (2), 186–208
- [71] Garg, A., Curtis, J. and Halper, H. (2003): Quantifying the financial impact of IT security breaches. *Information Management & Computer Security*, **11** (2), 74–83
- [72] Giovannetti, E., Neuhoff, K. and Spagnolo, G. (2005): Agglomeration in Internet co-operation peering agreements. *Cambridge Working Papers in Economics*, 0505. <http://econpapers.repec.org/paper/camcamdae/0505.htm> (last access: 10 Jan 2008)

- [73] Gordon, L. A., Loeb, M., Lucyshyn, W. (2003): Sharing information on computer systems security: An economic analysis. *J. Accounting Public Policy*, **22** (6), 461–485
- [74] Gordon, L. A. , Loeb, M., Sohail, T. (2003): A framework for using insurance for cyber-risk management. *Comm. ACM*, **46** (3), 81–85
- [75] Heise Online (2002): Tux takes its seat in Germany’s federal parliament. 28 February, at <http://www.heise.de/english/newsticker/news/25255>
- [76] House of Lords Science and Technology Committee (2007): *5th Report of Session 2006–07, Personal Internet Security*. <http://www.publications.parliament.uk/pa/ld200607/ldselect/ldsctech/165/165i.pdf> (last access: 22 Nov 2007)
- [77] Hovav, A. and D’Arcy, J. (2003): The impact of denial-of-service attack announcements on the market value of firms. *Risk Management and Insurance Review*, **6** (2), 97–121
- [78] Hovav, A. and D’Arcy, J. (2004): The impact of virus attack announcements on the market value of firms. *Information Systems Security*, **13** (2), 32–40
- [79] D’Ignazio, A. and Giovannetti, E. (2005): Spatial dispersion of peering clusters in the European Internet. *Cambridge Working Papers in Economics*, 0601. <http://econpapers.repec.org/paper/camcamdae/0601.htm>
- [80] D’Ignazio, A. and Giovannetti, E. (2006): ‘Unfair’ discrimination in two-sided Peering? Evidence from LINX, Cambridge Working Papers in Economics, 0621. <http://econpapers.repec.org/paper/camcamdae/0621.htm> (last access: 10 Jan 2008)
- [81] IBM (2007): Cyber attacks on the rise. IBM 2007 Midyear Report. http://www.iss.net/documents/whitepapers/x-force_threat_exec_brief.pdf (last access: 10 Jan 2008)
- [82] Insurance Information Institute (2003): Computer security-related insurance issues. <http://www.iii.org/media/hottopics/insurance/computer/> (historical state: 01 Oct 2004)
- [83] Interexport Management Systems (2006): Impact assessment study on possible options for the modification of the EU regime on export control of dual-use goods and technologies, 13 February, at <http://trade.ec.europa.eu/doclib/html/127589.htm> (last access: 10 Jan 2008)
- [84] International Organization for Standardization (2006): ISO/IEC 26300:2006 Information technology – Open Document Format for Office Applications (OpenDocument) v1.0.
- [85] Kannan, K. and Telang, R. (2005): Market for software vulnerabilities? Think Again. *Management Science*, **51** (5), 726–740

- [86] Keizer, G. (2007): Gartner: Bug bounty hunting is 'risky endeavour', *Computer-world* (2 May), at <http://www.computerworlduk.com/management/security/cybercrime/news/index.cfm?newsid=2836>
- [87] Keizer, G. (2007): Newest Windows update snafu puzzles Microsoft, *PC World* (16 October), at <http://www.pcworld.com/article/id,138495-pg,1/article.html>
- [88] Kesan, J. P., Majuca, R. P., Yurcik, W. (2005): Cyberinsurance as a market-based solution to the problem of cybersecurity – A case study. *Workshop in the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA. <http://www.infoecon.net/workshop/pdf/42.pdf> (last access: 3 Dec 2007)
- [89] Knuth, D. (2002): All questions answered. *Notices of the AMS*, **49** (3), 318–324 <http://www.ams.org/notices/200203/fea-knuth.pdf>
- [90] Kunreuther, H., Heal, G. (2003): Interdependent security. *Journal of Risk and Uncertainty*, **26** (2/3), 231–249
- [91] Lesk, M. (2007): The new front line: Estonia under cyberassault. *IEEE Security and Privacy*, **5** (4), 76–79
- [92] Majuca, R. P., Yurcik, W., Kesan, J. P. (2006): The evolution of cyberinsurance. *ACM Computing Research Repository (CoRR)*, Technical Report cs.CR/0601020. <http://uk.arxiv.org/ftp/cs/papers/0601/0601020.pdf> (last access: 30 Nov 2007)
- [93] Marcellis-Warin, N., Michel-Kerjan, E. (2001): The public-private sector risk-sharing in the French insurance 'Cat. Nat. System'. *CIRANO Séries Scientifique*, No. 2001s-60. <http://www.cirano.qc.ca/pdf/publication/2001s-60.pdf> (last access: 30 Nov 2007)
- [94] McAfee Inc. (2007): *McAfee Virtual Criminology Report*. http://www.mcafee.com/us/research/criminology_report/default.html (last access: 8 Dec 2007)
- [95] McClendon, M. J. (1991): Acquiescence and recency response-order effects in interview surveys. *Sociological Methods and Research*, **20**, 60–103
- [96] McPherson, D., Labovitz, C., Hollyman, M. (2007): *Worldwide Infrastructure Security Report*, vol. 3, Arbor Networks. <http://www.arbornetworks.com/report>
- [97] Manchester Evening News (2002): Blaze vandals sever Internet links, (23 October), at http://www.manchestereveningnews.co.uk/news/s/22/22480_blaze_vandals_sever_internet_links.html
- [98] Melander, I. (2007): Web search for bomb recipes should be blocked: EU. Reuters. <http://www.reuters.com/article/internetNews/idUSL1055133420070910>
- [99] Microsoft Corporation (2007): *Microsoft Security Intelligence Report (January – June 2007)*. <http://www.microsoft.com/downloads/details.aspx?FamilyId=4EDE2572-1D39-46EA-94C6-4851750A2CB0&displaylang=en> (last access: 10 Jan 2008)

- [100] Miller, C. (2007): The legitimate vulnerability market. *Workshop on the Economics of Information Security (WEIS)*, Carnegie Mellon Univ., Pittsburgh, PA. <http://weis2007.econinfosec.org/papers/29.pdf> (last access: 13 Nov 2007)
- [101] Moore, T., Clayton, R. (2007): Examining the impact of website take-down on phishing. *Proc. of Anti-Phishing Working Group eCrime Researcher's Summit (APWG eCrime)*, ACM Press, New York, 1–13
- [102] Mozilla Corporation (2007): Mozilla security bug bounty program. <http://www.mozilla.org/security/bug-bounty.html> (last access: 22 December 2007)
- [103] Mulligan, D. K., Bamberger, K. A. (2007): Security breach notification laws: Views from chief security officers. Samuelson Law, Technology & Public Policy Clinic, Univ. of California, Berkeley School of Law. http://www.law.berkeley.edu/clinics/samuelson/cso_study.pdf (last access: 7 Dec 2007)
- [104] Nader, R. (1965): *Unsafe at Any Speed*. Grossman Publishers Inc., New York.
- [105] Nardo, M. et al. (2005): Handbook on constructing composite indicators: Methodology and user guide. *OECD Statistical Working Paper*.
- [106] National Conference of State Legislatures (2007): Breach of information. <http://www.ncsl.org/programs/lis/cip/priv/breach.htm>
- [107] Office of Fair Trading (2003): Payment systems. http://www.oft.gov.uk/advice_and_resources/resource_base/market-studies/payment-systems
- [108] Ogut, H., Menon N., Rangunathan, S. (2005): Cyber insurance and IT security investment: Impact of independent risk. *Workshop in the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA. <http://www.infosecon.net/workshop/pdf/56.pdf> (last access: 3 Dec 2007)
- [109] OpenDNS (2007): OpenDNS shares April 2007 PhishTank statistics, Press Release, 1 May. http://www.opendns.com/about/press_release.php?id=14
- [110] Ozment, A. (2004): Bug auctions: Vulnerability markets reconsidered. *Workshop on the Economics of Information Security (WEIS)*, University of Minnesota, Minneapolis, MN. <http://www.dtc.umn.edu/weis2004/ozment.pdf> (last access: 13 Nov 2007)
- [111] Ozment, A. (2005): The likelihood of vulnerability rediscovery and the social utility of vulnerability hunting. *Workshop in the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA. <http://www.infosecon.net/workshop/pdf/10.pdf> (last access: 16 Nov 2007)
- [112] Ramzan, Z., Wüest, C. (2007): Phishing attacks: Analyzing trends in 2006. *Fourth Conference on Email and Anti-Spam*, Mountain View, CA. <http://www.ceas.cc/2007/papers/paper-34.pdf>

- [113] PITCOM (2006): Critical national infrastructure, briefings for parliamentarians on the politics of information technology. <http://www.pitcom.org.uk/briefings/PitComms1-CNI.doc> (last access: 10 Jan 2008)
- [114] Privacy Rights Clearinghouse (2005): A chronology of data breaches. <http://www.privacyrights.org/ar/ChronDataBreaches.htm>
- [115] Rice, D. (2007): *Geekonomics – The Real Cost of Insecure Software*. Addison-Wesley, New York.
- [116] Robinson, W. S. (1950): Ecological correlations and the behavior of individuals. *American Sociological Rev.* **15**, 351–357
- [117] Sachs, S.E. (2006): From St. Ives to cyberspace: The modern distortion of the medieval 'law merchant'. *American Univ. Int. Law Rev.*, **21** (5), 685–812. <http://ssrn.com/id=830265>
- [118] Sans Institute (2007): SANS Top-20 2007 security risks. <http://www.sans.org/top20/>
- [119] Schechter, S. E. (2004): *Computer Security Strength & Risk: A Quantitative Approach*. PhD thesis, Harvard University, Cambridge, MA
- [120] Schmidt, K. M., Schnitzer, M. (2003): Public subsidies for open source? Some economic policy issues of the software market. *CEPR Discussion Paper*, No. 3793
- [121] Schwalb, M. (2007): Exploit derivatives & national security. *Yale J. Law and Technology*, **9**, 162–192
- [122] Security Focus Inc. (2007): Bugtraq mailing list. <http://www.securityfocus.com/archive/1> (last access: 6 Dec 2007)
- [123] Serjantov, A. and Clayton, R. (2005): Modelling incentives for e-mail blocking strategies. *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA. <http://www.cl.cam.ac.uk/~rnc1/emailblocking.pdf> (last access: 10 Jan 2008)
- [124] Shapiro, C., Varian, H. R. (1999): *Information Rules. A Strategic Guide to the Network Economy*. Harvard Business School Press, Boston, MA
- [125] Shostack, A., Syverson, P. (2004): What price privacy? In J. Camp and S. Levis (eds.): *Economics of Information Security*, Kluwer Academic Publishers, Boston, 129–142
- [126] Silicon.fr (2007): l'Open Source pésera de plus en plus dans l'administration, (13 June). <http://www.silicon.fr/fr/news/2007/06/13/france-l-open-source-va-prendre>
- [127] St Albans District Council vs ICL (1996), at http://www.smaldonado.com/marcos/docs/it_case_su_uk_en.html

- [128] Sungard Availability Services (2006): 11 December 2005, Buncefield explosion, a Northgate Information Solutions case study. <http://www.availability.sungard.com/United+Kingdom/Resources/Case+Studies/Northgate+Information+Solutions.htm> (last access: 10 Jan 2008)
- [129] Sutton, M. and Nagle, F. (2006): Emerging economic models for vulnerability research. *Workshop on the Economics of Information Security*, Univ. of Cambridge, UK. <http://weis2006.econinfosec.org/docs/17.pdf> (last access: 13 Nov 2007)
- [130] Symantec (2007): *Internet Security Threat Report*. <http://www.symantec.com/business/theme.jsp?themeid=threatreport> (last access: 6 Dec 2007)
- [131] Telang, R. and Wattal, S. (2005): Impact of software vulnerability announcements on the market value of software vendors – an empirical investigation. *Workshop on the Economics of Information Security*, Harvard Univ., Cambridge, MA. http://infoecon.net/workshop/pdf/telang_wattal.pdf (last access: 13 Nov 2007)
- [132] United Kingdom (1990): Computer Misuse Act, c.18. <http://www.statutelaw.gov.uk/content.aspx?ActiveTextDocId=1353366>
- [133] United Kingdom Government (2007): The Government reply to the fifth report from the House of Lords Science and Technology Committee, Session 2006-07, HL Paper 165, Personal Internet Security. <http://www.official-documents.gov.uk/document/cm72/7234/7234.pdf> (last access: 22 Nov 2007)
- [134] United States Federal Trade Commission (2006): Zango, Inc. settles FTC charges. <http://www.ftc.gov/opa/2006/11/zango.shtm> (last access: 10 Jan 2008)
- [135] West, J. (2006): The economic realities of open standards: Black, white and many shades of gray. In S. Greenstein and V. Stango (eds.): *Standards and Public Policy*, Cambridge University Press, Cambridge
- [136] Wolfers, J. and Zitzewitz, E. (2004): Prediction markets. *J. Economic Perspectives*, **18** (2), 107–126
- [137] Zetter, K. (2005): Router flaw is a ticking bomb. *Wired* (1 August), at <http://www.wired.com/politics/security/news/2005/08/68365> (last access: 10 Jan 2008)
- [138] Zhuge, J., Holz, T., Han, X., Guo, J., Zou, W. (2007): Characterizing the IRC-based botnet phenomenon. Informatik Tech. Report TR-2007-010. <http://honeyblog.org/junkyard/reports/botnet-china-TR.pdf> (last access: 13 Dec 2007)
- [139] Zhuge, J., Holz, T., Song, C., Guo, J., Han, X., Zou, W. (2007): Studying malicious websites and the underground economy on the Chinese web. Informatik Tech. Report TR-2007-011. <http://honeyblog.org/junkyard/reports/www-china-TR.pdf> (last access: 13 Dec 2007)

A Information society indicators on security

Table 6: Information society indicators: security problems of individuals

Percentage of individuals who used Internet within the last year and have, in the last 12 months, experienced the following security problem ...									
	Fraudulent payment (credit or debit) card use			Abuse of personal information sent on the Internet			Computer virus res- ulting in loss of in- formation or time		
	2003	2004	2005	2003	2004	2005	2003	2004	2005
Bulgaria	–	0.3	–	–	3.0	–	–	37.8	–
Czech Republic	0.1	0.1	–	0.1	0.2	–	15.3	13.9	14.0
Denmark	0.8	1.1	1.2	0.4	1.1	1.5	27.6	30.1	35.0
Germany	–	–	–	4.3	2.7	2.0	13.1	35.0	33.3
Estonia	–	0.1	–	–	–	–	–	19.6	10.0
Ireland	0.7	1.1	0.6	2.4	1.8	1.3	11.6	24.8	16.6
Greece	0.1	0.1	0.4	1.2	0.8	0.5	14.7	12.0	17.9
Spain	–	0.8	1.7	–	18.5	15.4	–	50.8	47.8
Italy	–	–	0.7	–	–	4.0	–	–	41.3
Cyprus	–	0.9	0.5	–	4.0	8.9	–	27.0	24.5
Latvia	–	0.4	0.2	–	1.2	0.3	–	28.7	17.1
Lithuania	–	0.2	0.6	–	0.8	0.7	–	39.8	39.5
Luxembourg	1.5	0.6	1.4	4.1	9.8	6.3	24.9	49.8	46.0
Hungary	–	0.4	0.3	–	1.8	2.4	–	34.1	29.6
Netherlands	–	–	0.9	–	–	2.3	–	–	30.7
Austria	0.9	1.0	1.4	2.2	2.1	1.6	15.2	29.8	26.8
Poland	–	0.3	0.9	–	2.2	2.0	–	29.5	31.6
Portugal	–	–	–	4.2	1.4	2.0	14.0	17.5	23.4
Romania	–	0.1	–	–	0.4	–	–	5.2	–
Slovenia	–	0.7	–	–	1.4	–	–	33.9	39.8
Slovakia	–	0.3	0.2	–	2.8	1.0	–	29.2	25.9
Finland	0.2	0.0	–	3.6	4.5	2.8	13.1	26.6	31.0
Sweden	1.1	1.2	0.9	8.6	7.3	4.0	16.7	24.7	24.4
United Kingdom	1.7	2.4	3.3	3.2	3.3	3.1	26.6	29.8	37.4
EU 15	–	1.0	1.4	–	5.3	4.1	–	33.8	35.5
EU 27	–	0.9	1.3	–	4.6	3.8	–	32.3	34.4

Source: Eurostat. No data available for Belgium, France and Malta.

Table 7: Information society indicators: security problems of enterprises

Percentage of enterprises with Internet access having encountered the following security problem in the last 12 months ...									
	unauthorised access			blackmail & threats			virus attack		
	2003	2004	2005	2003	2004	2005	2003	2004	2005
Belgium	4	4	3	1	0	0	37	30	23
Bulgaria	–	3	2	–	0	0	–	25	19
Czech Republic	–	3	3	–	2	–	–	30	25
Denmark	4	4	5	1	–	–	46	32	24
Germany	–	2	1	–	1	0	–	24	21
Estonia	–	2	2	–	1	1	–	40	23
Ireland	–	4	4	–	2	2	–	45	39
Greece	5	3	3	0	1	0	49	31	27
Spain	–	3	3	–	0	0	–	33	27
Italy	4	2	4	0	1	1	53	24	50
Cyprus	–	2	1	–	0	0	–	33	35
Latvia	–	3	2	–	0	0	–	22	21
Lithuania	–	2	3	–	1	1	–	40	36
Luxembourg	3	3	3	1	1	1	28	36	22
Hungary	–	1	6	–	1	2	–	28	65
Malta	4	–	5	0	–	1	29	–	42
Netherlands	3	4	1	0	0	0	34	46	21
Austria	3	4	2	0	1	0	33	34	35
Poland	–	1	1	–	0	0	–	26	25
Portugal	4	4	–	1	1	–	28	34	13
Romania	–	2	–	–	0	–	–	29	–
Slovenia	–	2	2	–	1	0	–	31	31
Slovakia	–	1	1	–	0	0	–	22	20
Finland	4	4	4	–	0	–	41	53	55
Sweden	2	3	2	0	0	0	32	31	25
United Kingdom	–	–	3	–	–	0	–	–	22
EU 15	–	2	2	–	1	0	–	29	28
EU 27	–	2	2	–	1	0	–	29	29

Note: excluding financial sector

Source: Eurostat. No data available for France.

Table 8: Information society indicators: security problems of *large* enterprises

Percentage of large enterprises with Internet access having encountered the following security problem in the last 12 months ...									
	unauthorised access			blackmail & threats			virus attack		
	2003	2004	2005	2003	2004	2005	2003	2004	2005
Belgium	5	3	3	1	0	1	40	39	29
Bulgaria	–	2	2	–	1	0	–	30	27
Czech Republic	–	5	4	–	4	–	–	36	31
Denmark	7	–	5	1	–	–	49	46	36
Germany	–	3	2	–	0	0	–	41	33
Estonia	–	5	2	–	4	2	–	46	35
Ireland	–	7	2	–	1	1	–	55	52
Greece	3	3	3	2	1	1	54	34	25
Spain	–	3	3	–	0	0	–	48	37
Italy	5	3	5	0	1	1	66	41	63
Cyprus	–	0	0	–	0	0	–	36	30
Latvia	–	2	–	–	–	–	–	30	39
Lithuania	–	1	3	–	1	1	–	61	60
Luxembourg	3	7	3	0	0	0	30	32	12
Hungary	–	3	10	–	0	4	–	44	81
Malta	2	–	0	0	–	0	39	–	26
Netherlands	7	5	2	1	1	0	41	60	29
Austria	4	2	1	0	0	0	30	37	31
Poland	–	1	2	–	0	0	–	48	41
Portugal	4	5	2	1	2	1	41	44	22
Romania	–	3	–	–	0	–	–	44	–
Slovenia	–	3	2	–	1	0	–	51	45
Slovakia	–	1	1	–	0	0	–	22	26
Finland	7	4	4	–	–	–	48	63	63
Sweden	5	5	3	1	0	0	43	55	44
United Kingdom	–	–	3	–	–	0	–	–	29
EU 15	–	3	3	–	0	0	–	44	36
EU 27	–	3	3	–	1	1	–	43	37

Note: excluding financial sector. Enterprises with more than 250 employees

Source: Eurostat. No data available for France.

Table 9: Information society indicators: individuals' use of security technology

Percentage of individuals who used Internet within the last 3 months and have, in the last 3 months, ...						
	installed a virus checking program		updated a virus checking program		used online-authentication ^{a)}	
	2003	2004	2003	2004	2003	2004
Bulgaria	–	49	–	37	–	22
Czech Republic	–	–	–	–	–	30
Denmark	26	23	57	60	59	64
Germany	29	39	33	46	28	29
Estonia	–	1	–	1	–	1
Ireland	19	27	20	29	19	21
Greece	40	43	30	31	18	19
Cyprus	–	28	–	77	–	39
Latvia	–	31	–	31	–	23
Lithuania	–	18	–	18	–	20
Luxembourg	58	–	57	–	41	–
Hungary	–	55	–	46	–	28
Austria	30	34	31	42	23	28
Poland	–	48	–	35	–	25
Portugal	28	36	31	44	19	30
Romania	–	26	–	14	–	5
Slovenia	–	37	–	48	–	81
Slovakia	–	36	–	50	–	30
Finland	22	26	35	48	75	66
Sweden	26	25	40	48	56	51
United Kingdom	33	39	37	42	38	32
EU 27	–	–	–	–	–	32

Note: ^{a)} password, PIN, digital signature

Source: Eurostat.

No data available for BE, ES, FR, IT, MT and NL.

Table 10: Information society indicators: enterprises' use of security technology

Percentage of enterprises with Internet access having taken precautions.								
	all enterprises				large enterprises ^{a)}			
	2003	2004	2005	2006	2003	2004	2005	2006
Belgium	97	99	99	99	99	100	99	100
Bulgaria	–	100	–	97	–	100	–	99
Czech Republic	–	89	90	99	–	99	98	99
Denmark	96	96	98	99	99	100	100	100
Germany	–	98	98	99	–	99	100	98
Estonia	–	94	96	97	–	100	100	100
Ireland	52	93	98	98	71	99	100	99
Greece	96	97	98	99	100	100	99	100
Spain	43	99	99	99	77	100	100	100
France	–	–	–	96	–	–	–	100
Italy	100	99	99	100	100	100	100	100
Cyprus	–	99	99	99	–	100	100	100
Latvia	–	96	92	95	–	99	99	99
Lithuania	–	94	94	95	–	100	100	100
Luxembourg	96	97	98	99	97	97	100	99
Hungary	–	92	93	97	–	96	99	100
Malta	79	–	98	–	76	–	100	–
Netherlands	97	95	98	99	100	99	100	100
Austria	97	98	99	99	99	100	100	100
Poland	–	93	91	94	–	100	99	100
Portugal	92	95	96	98	98	100	100	100
Romania	–	66	–	93	–	84	–	98
Slovenia	–	98	99	99	–	100	100	100
Slovakia	–	97	96	97	–	100	100	99
Finland	96	97	99	100	99	100	99	100
Sweden	97	98	99	99	99	–	99	100
United Kingdom	–	–	98	99	–	–	100	100
EU 15	–	98	98	99	–	100	100	99
EU 27	–	97	97	98	–	99	100	99

Note: excluding financial sector. ^{a)} Enterprises with more than 250 employees
Source: Eurostat.

Table 11: Information society indicators: enterprises' update practices

Percentage of enterprises with Internet access that have installed security devices on their PCs and updated them within the last three months						
	all enterprises			large enterprises ^{a)}		
	2003	2004	2005	2003	2004	2005
Belgium	86	91	85	97	98	95
Bulgaria	–	77	59	–	86	78
Czech Republic	–	74	77	–	93	94
Denmark	82	90	88	96	98	97
Germany	76	86	69	92	98	85
Estonia	–	57	66	–	82	90
Ireland	–	78	79	–	93	98
Greece	72	70	78	90	88	90
Spain	34	84	83	71	96	95
Italy	63	83	87	86	96	97
Cyprus	–	81	83	–	95	96
Latvia	–	72	64	–	91	90
Lithuania	–	56	67	–	86	91
Luxembourg	77	87	88	88	95	98
Hungary	–	61	72	–	71	83
Malta	66	–	90	71	–	100
Netherlands	86	86	72	95	95	88
Austria	88	93	90	98	99	98
Poland	–	58	68	–	88	94
Portugal	68	79	80	91	96	95
Romania	–	43	–	–	61	–
Slovenia	–	70	78	–	90	93
Slovakia	–	84	83	–	99	94
Finland	86	92	91	95	97	95
Sweden	83	93	93	96	99	98
United Kingdom	–	–	68	–	–	84
EU 15	68	85	77	89	97	89
EU 27	–	81	77	–	93	89

Note: excluding financial sector. ^{a)} more than 250 employees

Source: Eurostat. No data available for France.

Table 12: Information society indicators: perceived barriers to e-commerce

Percentage of individuals who, in the last 12 months, haven't ordered goods or services over the Internet, because of ...						
	security concerns		privacy concerns		security or privacy concerns	
	2004	2005	2004	2005	2005	2006
Belgium	–	–	–	–	–	26
Bulgaria	18	–	–	–	–	9
Czech Republic	–	–	–	10	–	8
Denmark	34	–	3	–	–	29
Germany	33	28	28	23	32	57
Estonia	–	11	–	7	12	21
Ireland	6	7	2	2	8	9
Greece	37	19	33	24	32	45
Spain	26	68	22	60	–	71
France	–	–	–	–	–	48
Italy	–	25	–	15	30	29
Cyprus	82	67	84	59	68	64
Latvia	6	7	3	4	9	7
Lithuania	5	12	–	10	15	18
Luxembourg	27	46	15	19	50	46
Hungary	–	35	–	30	38	36
Netherlands	–	38	–	30	42	41
Austria	–	20	–	15	22	24
Poland	31	7	34	6	9	10
Portugal	39	43	39	42	47	40
Romania	–	–	–	–	–	6
Slovenia	35	57	29	51	60	43
Slovakia	9	11	7	10	15	17
Finland	–	62	–	61	66	69
Sweden	–	14	–	11	16	19
United Kingdom	17	36	–	24	39	35
EU 15	–	37	–	30	34	45
EU 27	–	32	–	25	30	38

Source: Eurostat. No data available for Malta.

B Internet exchange points

Table 13: Internet exchange point sizes (number of participants)

Name	Full Name	Country	Size	Share
VIX	Vienna Internet Exchange	AT	39	100 %
BNIX	Belgian National Internet eXchange	BE	27	51 %
FreeBiX	FreeBiX	BE	26	49 %
NIX CZ	Neutral Internet Exchange	CZ	23	100 %
DE-CIX	Deutscher Commercial Internet Exchange	DE	209	64 %
KleyReX	KleyReX	DE	43	13 %
INXS	Internet Exchange Service	DE	17	5 %
B-CIX	Berlin Commercial Internet Exchange	DE	15	5 %
ECIX Duesseldorf	European Commercial Exchange Duesseldorf	DE	14	4 %
WORKIX	WORKIX Hamburg	DE	11	3 %
N-IX	Nuernberger Internet Exchange	DE	7	2 %
ECIX Berlin	European Commercial Exchange Berlin	DE	5	2 %
OCIX Duesseldorf	OpenCarrier e.G. Member IX Duesseldorf	DE	4	1 %
ECIX Leipzig	European Commercial Exchange Leipzig	DE	0	0 %
S-IX	Stuttgarter internet eXchange	DE	0	0 %
DIX	Danish Internet Exchange	DK	20	100 %
TIX-LAN	Tallinn Internet eXchange	EE	4	67 %
TLLIX	Tallinn Internet Exchange	EE	2	33 %
ESPANIX	Espana Internet Exchange	ES	20	59 %
CATNIX	Catalunya Neutral Internet Exchange	ES	10	29 %
Terremark NAP	Terremark NAP de las madrid	ES	3	9 %
MAD-IX	Madrid Internet Exchange	ES	1	3 %
EuskoNIX	Basque Country Internet Exchange Point	ES	0	0 %
FICIX	Finnish Communication and Internet Exchange	FI	14	100 %
PaNAP	Paris free NAP	FR	65	29 %
SFINX	Service for French INternet eXchange	FR	54	24 %
FreeIX	Free Internet eXchange	FR	52	23 %
PARIX	PARIX	FR	28	12 %
Pouix	Paris Operators for Universal Internet eXchange	FR	16	7 %
FNIX6	French National Internet Exchange IPv6	FR	3	1 %
Lyonix	Lyonix & the Lyon IX	FR	2	1 %
GEIX	Gigabit European Internet Exchange	FR	4	2 %
EuroGIX	EuroGIX	FR	0	0 %
MA-IX	Marseille Internet eXchange	FR	1	0 %
PhibIX	PhibIX Gix & Nap	FR	1	0 %
AIX	Athens Internet Exchange	GR	4	100 %
BiX	Budapest Internet Exchange	HU	7	100 %
INEX	Internet Neutral EXchange	IE	17	100 %
MIX-IT	Milano Internet eXchange	IT	29	67 %
NaMeX	Nautilus Mediterranean Exchange Point	IT	8	19 %
TOP-IX	Consorzio Top-IX	IT	6	14 %

Continued on next page

Name	Full Name	Country	Size	Share
LIX	Luxemburg Internet Exchange	LU	3	100 %
MIX Malta	Malta Internet Exchange	MT	0	100 %
AMS-IX	Amsterdam Internet Exchange	NL	239	76 %
NL-IX	Netherlands Internet Exchange	NL	63	20 %
GN-IX	Groningen Internet Exchange	NL	10	3 %
NDIX	Nederlands Duitse Internet Exchange	NL	3	1 %
WIX	Warsaw Internet eXchange	PL	4	50 %
PL-IX	Polish Internet eXchange	PL	4	50 %
GIGAPIX	GIGAbit Portuguese Internet eXchange	PT	13	100 %
BUHIX	Internet Exchange Bucharest	RO	0	100 %
NetNod Stockholm	NetNod Internet Exchange i Sverige AB	SE	61	66 %
SOL-IX	SOL-IX	SE	10	11 %
STHIX	Stockholm Internet exchange	SE	7	8 %
Netnod Malmoe	Netnod Internet Exchange Ab	SE	5	5 %
Netnod Sundsvall	Netnod Internet Exchange Ab	SE	5	5 %
Netnod Gothenburg	Netnod Internet Exchange Ab	SE	4	4 %
GIX	Gothenburg Internet Exchange	SE	0	0 %
SIX.SK	Slovak Internet Exchange	SK	5	100 %
LINX	London Internet Exchange Ltd.	UK	394	59 %
LONAP	London Network Access Point	UK	86	13 %
LIPEX	London Internet Providers Exchange	UK	64	10 %
XchangePoint	XchangePoint Europe	UK	40	6 %
MaNAP	Manchester Network Access Point	UK	26	4 %
XchangePoint Lon.	XchangePoint Europe	UK	14	2 %
MerieX	Meridian Gate Internet Exchange	UK	11	2 %
RBIEX	RBIEX Limited	UK	9	1 %
UK6X	BTexact Technologies IPv6	UK	8	1 %
MCIX	Manchester Commercial Internet Exchange	UK	7	1 %
ENLIX	Enlightened Internet Exchange	UK	4	1 %
eXpress	PacketExchange - eXpress	UK	2	0 %
UNION IXP	UNION IXP	UK	0	0 %
EarthNAP	EarthNAP	UK	0	0 %

Source: <http://www.peeringdb.com>

C Methodology

The work was carried out between October – December 2007 by the four authors with the bulk of the work done in November and December. During the research phase of the report we studied the literature on security economics and performed an analysis in which we considered the incentives of various stakeholders for improving security using the following framework.

1. Core platforms
 - Operating systems – Microsoft, Apple, Symbian, and free platforms such as Linux and OpenBSD
 - Communications systems – routers (Cisco) and phone switchgear
 - Server platforms – Microsoft, Apache
2. Core services
 - Internet service provision
 - Banking and payment services
 - Search, advertising, and other business-critical infrastructure
3. Services dependent on e-communication
 - E-government providers, from tax to company registration
 - Firms selling services online such as insurance, hotel bookings, entertainment and transport tickets
 - Firms selling intangible goods such as music downloads
4. Services becoming dependant on e-communication
 - Health care providers
 - Manufacturing automation
 - Supply chain management
 - Customer services (Call centres, CRM systems)
5. Security ‘providers’
 - Policy makers, public sector agencies (both civilian and defence), NGOs
 - Universities and other research institutions
 - Software developers (commercial and other)
 - Service firms (auditors, consultants)
6. Consumers

Many of these stakeholders are both producers and consumers of security. We therefore conducted a matrix analysis to tease out the incentive failures at each interface. For example, under the heading of ‘information asymmetries’, we looked at whether platform vendors, firms, other service providers and security providers were likely to improve measurability (for example, if it could differentiate themselves from other providers). We also looked at which providers have an incentive to inflate either threats or countermeasures, with negative consequences.

We then spoke to senior managers at a number of key companies, attended several conferences and progressively tested our emerging conclusions in discussion with expert colleagues. We would also like to acknowledge valuable input from a project being run in parallel with ours by the OECD, in which Professor Michael van Eeten and colleagues studied the economics of malware and conducted extensive structured interviews with stakeholders in order to determine incentives and attitudes [41].

This initial research led to a workshop in Brussels on December 10th at which draft conclusions and options were discussed with stakeholders (the ENISA Permanent Stakeholders’ Group, European officials, industry representatives and NGOs). Draft recommendations were presented at this meeting at which there was extensive feedback from stakeholders.

Our report was evolved throughout this process and then settled in late December 2007.