

Security Economics and European Policy

Ross Anderson, Rainer Böhme, Richard Clayton and Tyler Moore

March 1, 2008

1 Summary

In September 2007, we were awarded a contract by the European Network and Information Security Agency (ENISA) to investigate failures in the market for secure electronic communications within the European Union, and come up with policy recommendations. In the process, we spoke to a large number of stakeholders, and held a consultative meeting in December 2007 in Brussels to present draft proposals. This established that almost all of our proposals have wide stakeholder support. The formal outcome of our work was a detailed report, *'Security Economics and the Internal Market'*, that is due to be published by ENISA. This paper is a much abridged version (about half the length): in it, we present the recommendations we made, and then a summary of our reasoning. By way of disclaimer, we state that these recommendations are our own and do not necessarily reflect the policy of ENISA or any other European institution.

The background should be familiar enough. The direct cost to Europe of electronic crime, including both losses and protective measures, is measured in billions of euros; and growing public concerns about information security hinder the development of both markets and public services, causing 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. Privacy concerns are stalling the development of e-health and other systems.

Information security is now a mainstream political issue. An appropriate regulatory framework, which is just as important for protecting economic and other activity online as it is offline. The European Union already has a number of laws on matters from e-commerce through telecomms regulation to consumer protection and product liability that regulate online activity, but the pace of change has left a number of gaps. To close these, we make the following recommendations.

1.1 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 electronic 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 stakeholders 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.

2 Background

Since about 2004, volume crime has arrived on the Internet. 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 bank employee at a call center collecting password data for use by an accomplice. Such crimes were local and inefficient.

2.1 The villains’ PIN factory

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 on to third countries. 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

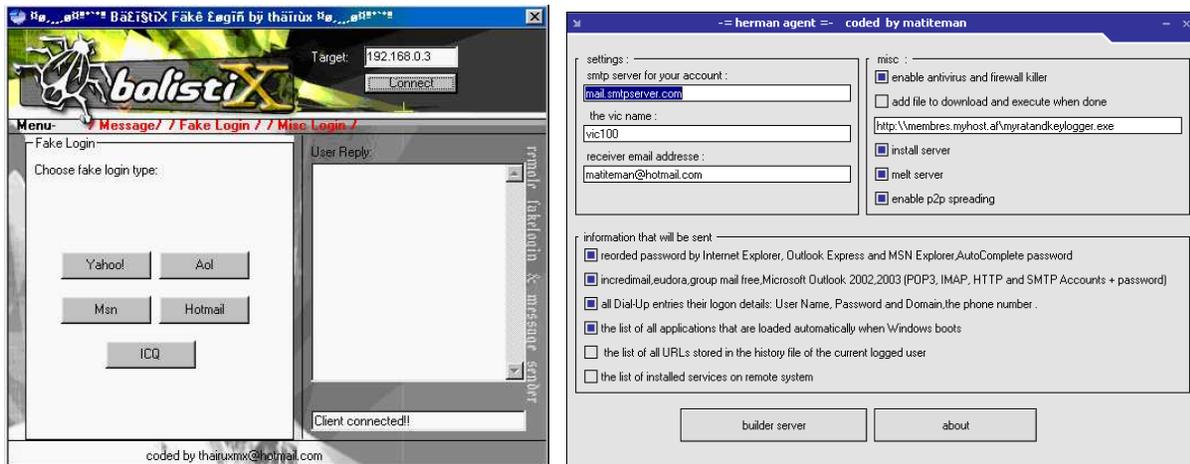


Figure 1: Web trojan generator interface (left) and data theft crimeware interface (right). Source: [36]

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 Explorer 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 [105]. 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 [36].

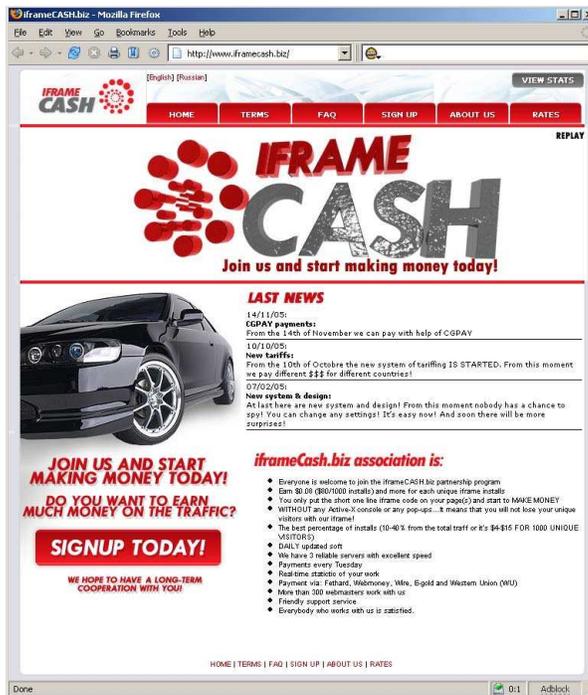


Figure 2: Crimeware affiliate marketing. Source: [36]

2.2 Security economics

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 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 [38].) It thus appears that the indirect costs of Internet insecurity are billions of euros more.

The economic study of information security 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.

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.

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.

2.3 The structure of our report

We used five general headings to classify and analyse the economic barriers to network and information security: information asymmetries, externalities, liability, diversity, and the fragmentation of legislation and law enforcement.

Information Asymmetries Asymmetric information can be a strong impediment to effective security. Akerlof's model of the 'market for lemons' [2] appears to apply to many security product markets. 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.

It has also long been known that we simply do not have good statistics on online crime, attacks and vulnerabilities. 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.

Externalities 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. As a result, a user who connects an unpatched computer to the Internet does not face the full economic consequences of her action.

A further set of externalities affect ISPs. 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 [98]) while large ISPs at present enjoy a certain impunity.

Network externalities also affect many protective measures. 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. In other circumstances, investments can be strategic complements: an individual taking protective measures may also protect others, inviting them to free-ride.

Liability dumping Firms seeking to manage risk often dump it on less powerful suppliers or customers. Software and service suppliers impose licenses on customers disclaiming all liability, including for security failures, and may also take 'consent' to the installation of spyware. This may delay the emergence of a market for more secure languages and tools, and lessen demand for the employment of professional software engineering methods.

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 [61], which recommended that the traditional consumer protection enshrined in banking law since the nineteenth century should be extended to electronic transactions too.

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 competition authorities; lack of diversity makes successful attacks more devastating – and harder to insure against, as high loss correlation renders some 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.

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 jurisdictions hinders rapid response. 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 send hot money through the banks of Member States with a relaxed attitude to asset recovery. This issue spills over to money laundering.

A serious problem is that traditional mechanisms for international police cooperation are too slow and expensive for the Internet age. They evolved when international investigations were infrequent and dealt with matters that were either procedurally simple (such as the extradition of a fugitive) or a large investigation of mutual interest (such as drug smuggling). They cannot cope well (or in some cases at all) with volume crime that crosses national boundaries.

3 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 3.1. This 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 3.2, we discuss other available data sources and requirements for robust security statistical indicators. In Section 3.3, we outline conditions for stakeholders to share relevant data and make recommendations to increase data-sharing.

3.1 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 [15], which came into force as Cal. Civil Code §1798.29, in July 2003. It requires public and private entities that conduct business in California to notify individuals whose personal data under is believed to have been acquired by an unauthorised person. This law was intended to give individuals the opportunity to protect their interests – such as putting a ‘lock’ on their file at credit agencies – and to motivate companies holding personal data to keep it secure in the first place. In some cases, such as the ChoicePoint scandal where criminals were able to access 163,000 credit reports, disclosure has had a substantial impact on the stock price [1] – not least because the regulator subsequently fined the company USD 15 million [51].

The California law has been followed by further laws in at least 34 other states [86], although they differ somewhat in their details. The variations have led to calls for a federal statute, but although bills have been introduced in Congress, none have had much success so far.

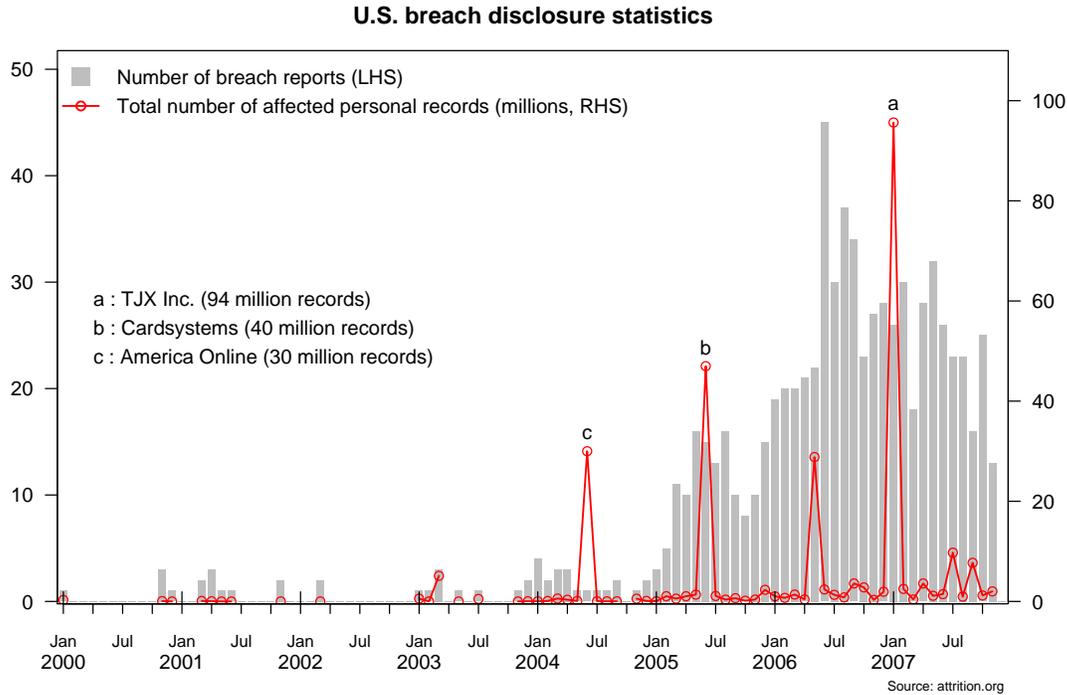


Figure 3: Bulge of breach reports after the introduction of disclosure laws in the US

The Privacy Rights Clearinghouse publishes a database of known security breaches and gives brief details of each one [92]. The number of records compromised now exceeds 215 million. Several research groups, above all the contributors to attrition.org¹, 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 entity [84]. The median, as a robust measure, is a moderate 8,000 records per breach.

Figure 4 shows the annual breakdown of breach disclosures by sector 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 (up to end November). Fraud and social engineering are 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 re-

¹<http://attrition.org>

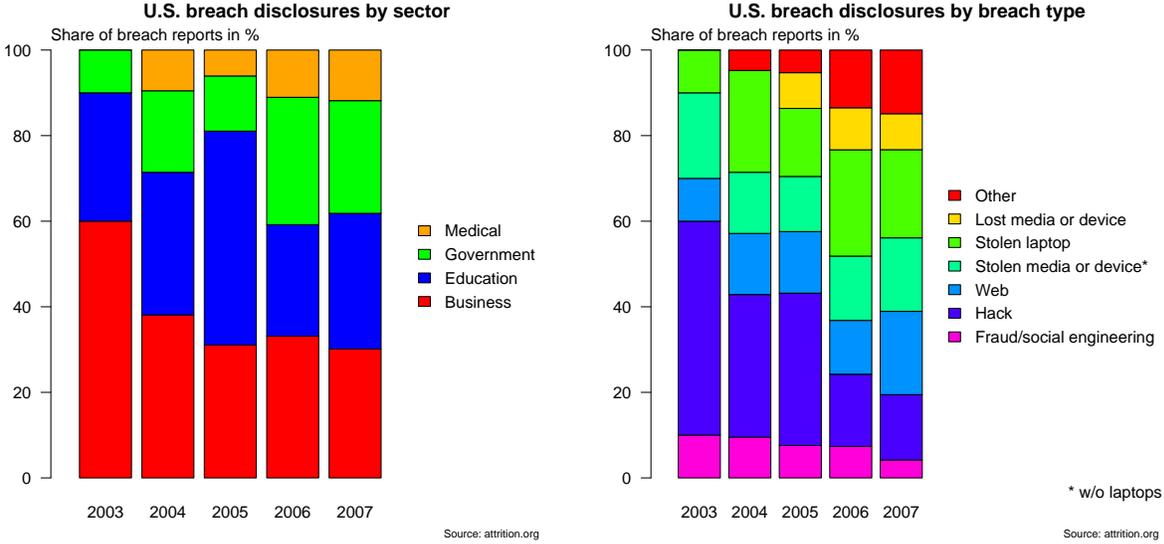


Figure 4: Distribution of breach reports across sectors (left) and breach types (right)

records. We computed the logarithm before calculating sector and type averages to account for the great variation in the number of records disclosed. Data losses are increasingly caused by accidents; hacks account for a diminishing, but still substantial, proportion, while breaches via the web compromise the fewest records. Businesses tend to put most records at risk, while the education sector exposes the fewest.

In Europe, a security breach notification law has been proposed that would require notification to be made where a network security breach was responsible for the disclosure of personal data [39]. This is a very narrow definition 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.

In the UK, the House of Lords Personal Internet Security inquiry [61] recommended that the UK bring in such a security breach notification law; the Government’s response was negative [107]; and the Lords are planning further hearings at the time of writing.

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 rather than just telecoms companies. Indeed, 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 protection. It does not impose the burden of a strict liability regime across the whole economy, but relies on ‘naming and shaming’. Competent firms 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.

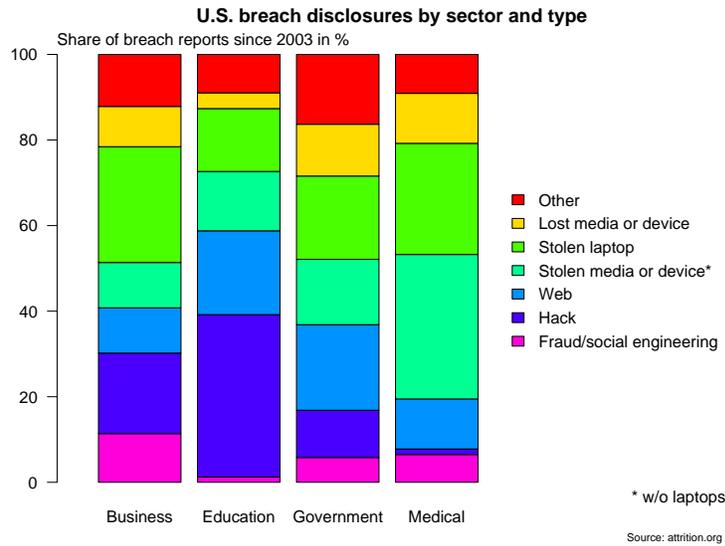


Figure 5: Breakdown by sector *and* breach type: Education is primarily hit by hacks while theft dominates in the medical sector

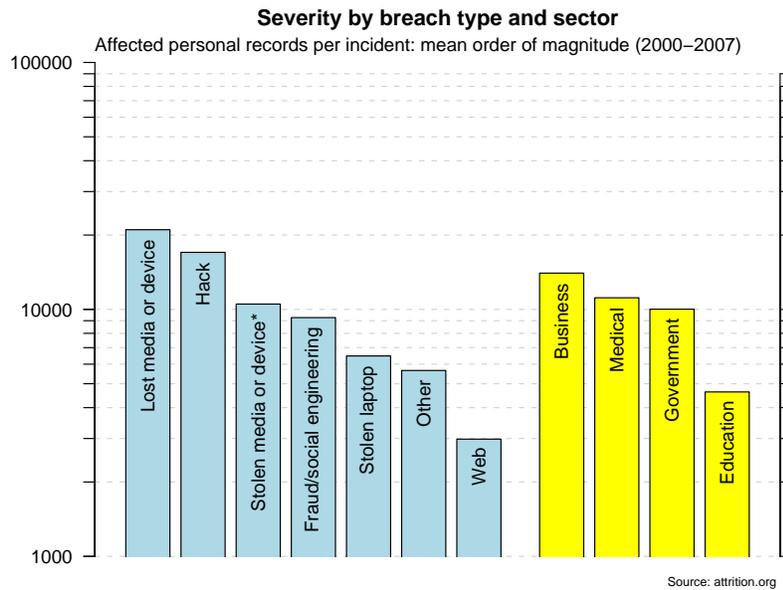


Figure 6: Log average of number of affected personal records per breach report broken down by breach type (left bars) and sector (right bars). Note the log scale.

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.

3.2 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 [54]. 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.

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.

Statistical data have many uses in the information security economy. They can mitigate information asymmetries by generating useful signals for economic decision making, whether by policymakers, firms or individuals. Data also help security professionals to plan and implement appropriate protection. Another key use is policy formation; yet another is academic research. Different users often have different requirements in terms of timeliness and resolution.

Sometimes the uses overlap. For example, Moore and Clayton have studied the effectiveness of phishing website removal countermeasures instigated by the banks and specialist ‘take-down’ companies [82]. They have found 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!

3.2.1 What statistics are already being collected?

ENISA has recently published a report that outlines over 100 sources of data on information security [17]. Published data comes in many forms.

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 [23]. We examine data from the CSI survey in Section 3.2.2.

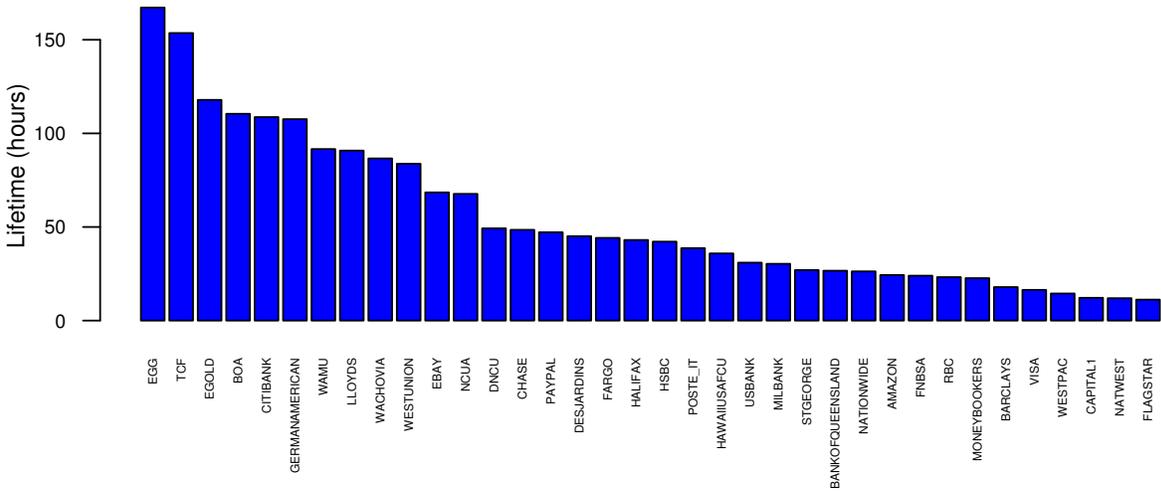


Figure 7: Phishing-site lifetimes per bank. Source: [82]

There has so far been one relevant Community initiative: in 2003, Eurostat started collecting data on Internet security issues from both individuals and enterprises in its ‘Community Surveys on ICT Usage’ [38].

Security *breach-disclosure reports* provide another useful data source. Groups such as attrition.org collate reports, which we discuss in Section 3.1 above.

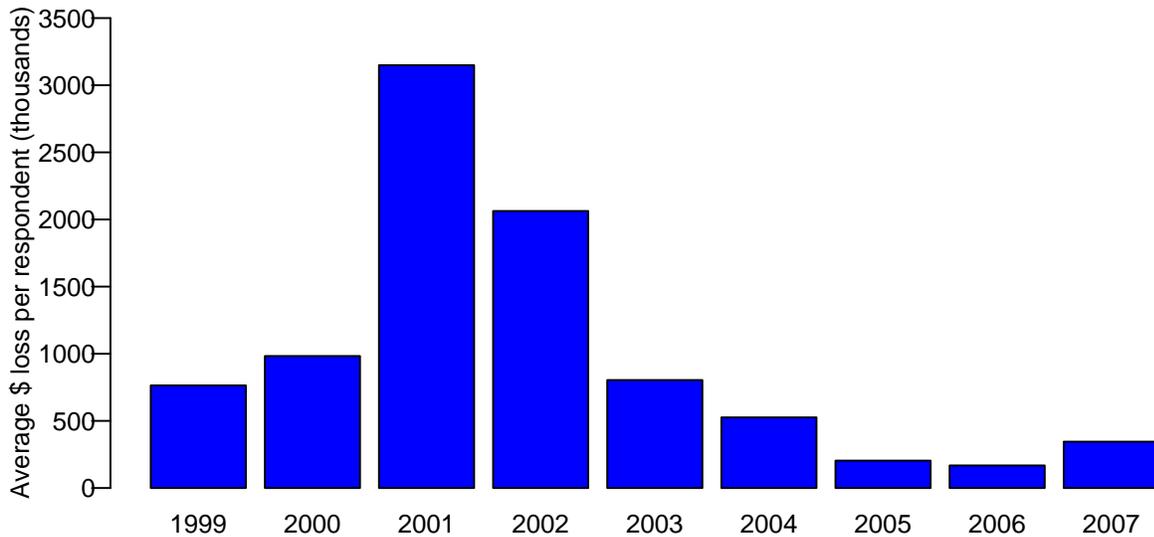
Direct observation is a third category of data collection on information security. Many security vendors regularly publish reports. For example, Symantec has published a semi-annual Internet Security Threat Report (ISTR) since 2002 [105]. Symantec directly measures many types of malicious activity using its global infrastructure of 40,000 sensors. We discuss Symantec’s report in Sections 3.2.2 and 5.5. Other security organisations such as McAfee [77], SANS [94], IBM [66] and Microsoft [80] 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 3.2.2 and 3.3.2) and APACS, the UK payments association (see discussion in Section 3.3.2). Finally, some academics conduct useful data collection and analysis. In this report, we refer to analysis of phishing websites by Moore and Clayton [82] in Figure 7 and malware tracking by Zhuge et al. [110] in Sections 5.5 and 7.2.

3.2.2 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.

The CSI survey² 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 losses due to various attack types.

²This was called the CSI/FBI Computer Crime and Security Survey until 2007.



Source: CSI 2007 Computer Crime and Security Survey

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

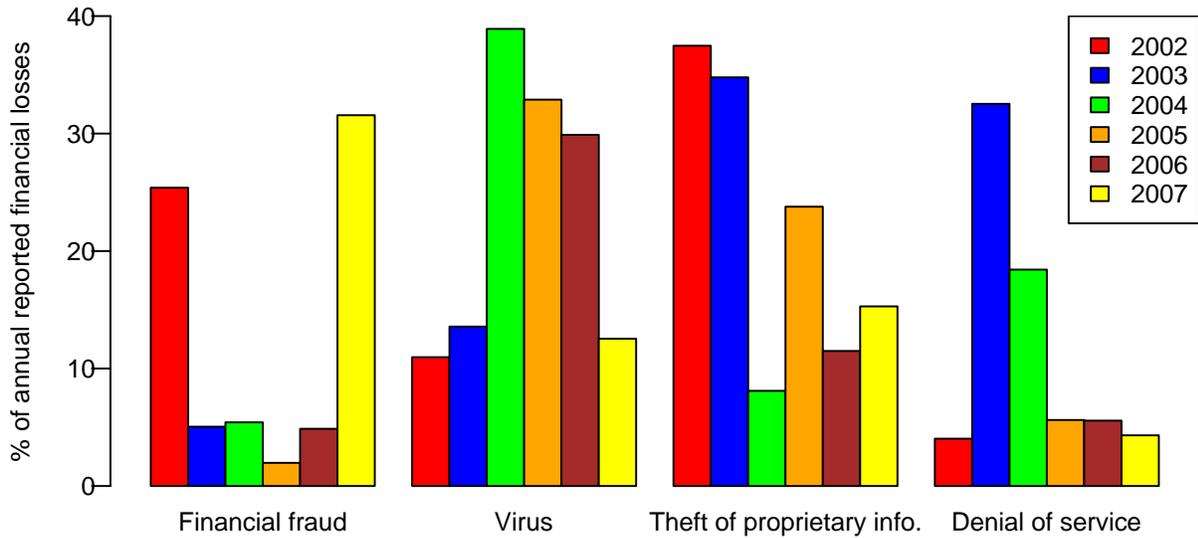
While there are undoubtedly problems associated with asking firms to self-report losses, there is no better measure of monetary losses at present.

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.

The Eurostat survey 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 problems, taken precautions, and installed security devices on their PCs. Individuals are further broken down by age group and residence, while data on enterprises are available for different firm sizes and sectors (excluding the financial sector).

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



Source: CSI 2002–2007 Computer Crime and Security Surveys

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

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.

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 a conflict of interest: their reports are published principally for marketing reasons. It is not surprising, then, that sometimes their data are consistently over-reported. For example, we studied more closely one statistic measured in several reports, which tracks the proportion of malicious code that exploits confidential

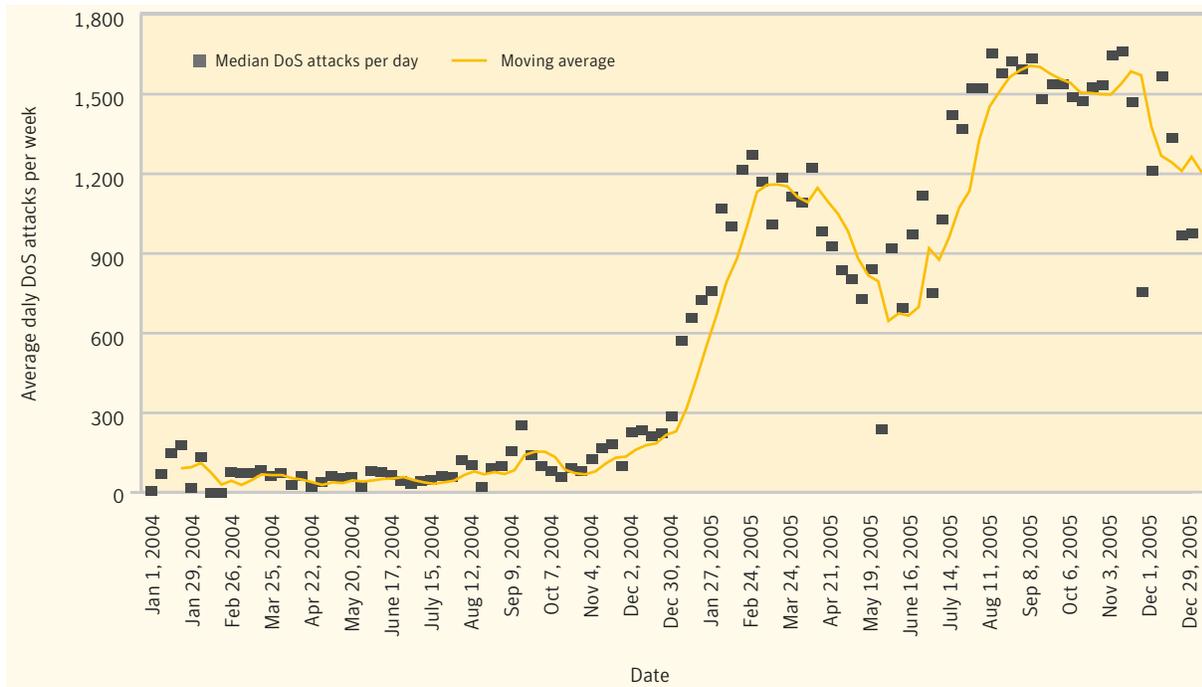


Figure 10: Denial-of-service attacks per week in 2004 and 2005 (Reproduced from Symantec Internet Security Threat Report IX)

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 is 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’.

Also, 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; domestic respondents are mostly unable to comprehend the questions [38], where for some countries the percentage of households who claim that they have updated their virus checking program in the last three months exceeds the percentage of household who report that they have installed one. Nearly all enterprises reported that they have taken precautions. Eurostat’s work on these indicators appears stalled, with the dim

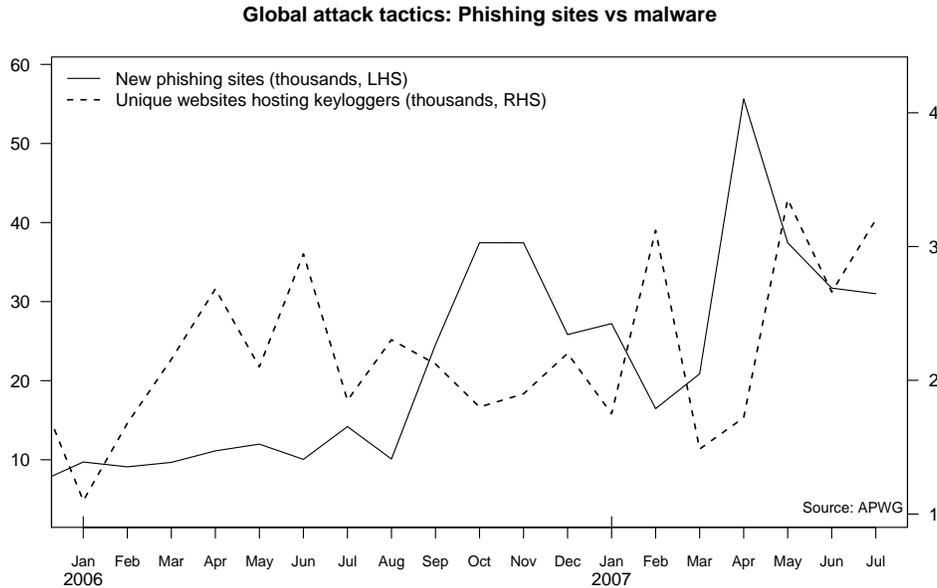


Figure 11: Attack trends vary rapidly over time. Phishing sites and keyloggers are substitutes, the Pearson correlation of first difference series is -0.4 .

outlook of a special module on security, scheduled for 2010 – at the last possible date of the current i2010 agenda [38].

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 crimefighters and service providers.

Robustness to short-lived tactics can be achieved by measuring losses. 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. Banks are likely to be a better reference source than ISPs. Unfortunately, only the UK banking industry provides dependable public statistics of electronic fraud losses.

3.2.3 Metrics derived from market price information

The efficient market hypothesis, in its strong form, suggests that stock market prices aggregate all information relevant for forming expectations of future profits [50]. If information security matters for a company, then the stock price 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 [49, 62], security incidents generally [16, 56], computer virus contagion [63], vulnerability disclosure in software products

(though smaller in magnitude) [106], and privacy breaches [1]. Another study also finds positive market reactions for listed security firms when news on security breaches is in the media [19]. Other co-variables 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).

These studies offer evidence that information security news impacts stock prices. However, their usefulness as a direct metric for security is limited. First, event studies capture short-term losses only, and second, event studies are limited to extreme events, such as attacks and security breaches. In general, stock prices aggregate too much diverse information. 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.

Many ideas have been put forward to use markets to extract security-related information – not just to align incentives, but also to provide new security metrics. A recent proposal [53] to track ‘underground market indices’ in IRC channels to feed forecasting and threat prediction tools may sound a bit like ‘the street price of drugs’, used in narcotics enforcement: these markets operate as exchange platforms for stolen credit card and identity details, hacked accounts, spam distribution and related services.

There has been more work on ‘vulnerability markets’ [11, 104, 81]: black markets [81], vulnerability brokers [104, 68], bug bounties [72, 83, 95, 69], and bug auctions [90, 11]. These all suffer from the usual problems of dealing with information goods. There are also suggested innovations using indirect markets, such as exploit derivatives [11, 108, 96]. Meanwhile, the price of cyber-insurance provides an indirect market measure of overall systems vulnerability [10]. (We will discuss cyber-insurance in much greater detail in Section 8.1.)

3.3 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.

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 3.2).

3.3.1 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 [35]. Managers may also be loth to disclose breaches in case they get fired [55]. Data from the 2007 CSI Computer Crime and Security Survey [23] report that negative publicity is the most-cited reason to abstain from

reporting a security problem (26 % of respondents name it). A breach-disclosure law, which we advocate, will blunt these disincentives.

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 [58]. Even the threat of such free-riding can stymie sharing. 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. But there is little evidence that this is a major concern in practice.

There can also be positive economic incentives for sharing security information. Gal-Or and Ghose developed a model of where sharing can work [55]: they argue that information sharing can encourage additional security investment. 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.

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; phishing statistics often seem particularly phishy. For example, the anti-phishing group PhishTank has boasted about the large number of sites it identifies [89], 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) [6].

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.

In more mature sectors of the economy, we can see useful examples of statistical institutions collecting business data jointly with industry bodies. For example, safety and accident statistics for cars are collected by police and insurers, while media circulation figures are typically collected by private firms, some of them jointly owned and controlled by publishers and advertisers.

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 [17]. 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. Hence mandatory reporting of particular indicators may be required for sharing to happen. Let us look now at the options.

3.3.2 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). 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 [27]. 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.

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 [5]; 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, 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 wiki³, categorise reports and facilitate understanding of reports without revealing details. 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 likely to be necessary.

Option 4: Aggregated fraud figures driven by ENISA Elsewhere in the financial industry, the incentives against sharing security-related information could be overcome with additional regulatory encouragement.

Individual banks are usually keen to keep such data private. But one notable exception is the behaviour of APACS, the UK payments association, which has published aggregated figures for the annual amount lost to phishing attacks, as well as ATM crime and other financial fraud [6]. While the incentives are against individual financial institutions

³<http://wiki.enisa.europa.eu>

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.

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, 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.

3.4 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. 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).

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'.

4 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.

4.1 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 smartcard industry is thus viewed as the poster case of economic development in the technology field being spurred by state intervention.

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. 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.

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.

4.2 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.

4.2.1 Control points

There are a number of *control points* where we might possibly do something about system insecurity. In Section 5 we will discuss what can be done with the stick of liability. 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.

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.

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 [78]. 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.

4.2.2 Policy options for coping with externalities

So if ISPs should take actions to raise the level of end-user security, then how can we best 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.

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. We already recommended that ENISA collect and publish data about this (Recommendation 3). But decent statistics alone are unlikely to be enough. For many years, governments published car-theft statistics and left it to the car industry to make their cars harder to steal. Yet 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.

So 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. But there are several reasons to doubt whether self-regulation will work any better with ISPs than it did with the auto industry. First, it has not worked very well so far. An OECD report investigating the economics of malware [35] argues that some positive incentives exist for ISPs to take precautionary measures, such as the high cost of customer support. But this doesn't work when the malware is designed to be undetectable, as is increasingly the case.

Some ISPs have taken action, but the poor performance of other ISPs overshadows this. Overall Internet security is down to the 'weakest link': attackers identify and exploit the worst-performing ISPs. If one ISP cleans up infected machines, the attackers may simply compromise more machines elsewhere. 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 pressure to clean up their act.

Option 2: Taxing 'digital pollution' ISPs that fail to take suitable measures to prevent their customers becoming a nuisance to other Internet users could be taxed directly. For example, a public authority might apply a penalty charge to ISPs that fail to take down infected machines within a fixed period, or levy an annual tax based on statistical measurements of emitted malware. Taxation is likely to be vehemently resisted by ISPs; a stakeholder at our consultative meeting inveighed against 'punishing the innocent'. Taxing customers directly would be even more controversial. There is also the concern that a heavily-policed takedown regime might be vulnerable in some Member States to capture by vested interests – for example, music companies wanting the government to be more vigorous in disconnecting the PCs of children accused of downloading music.

Option 3: A cap-and-trade system One suggested alternative is a cap-and-trade system, as used already with carbon credits. 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 can 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: the potentially high transaction cost of lawsuits; and 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. Upon notification of malicious activity, ISPs should place the machine into quarantine, clean up the offending content and reconnect the user as soon as possible. Unfortunately, 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 cleanup. This process could be formalised into a quarantine 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.

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. 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. So there has to be a ‘put-back’ mechanism that users can invoke to get their ISPs to reconnect an incorrectly classified machine quickly by assuming liability for any wicked emanations. Consumers only need assume liability if they skip the quarantine process. In practice, we anticipate most consumers will elect to participate in the ISP’s cleanup service.

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? Do we need policemen in each ISP dealing with infected machines, or could the ISPs' own staff do it more efficiently and cheaply?

5 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.

5.1 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* [85], and by US case law enabling accident victims to sue the manufacturer and not just the driver or the car dealer [93]. This led to a change of attitude by car makers, helped along by a multitude of regulatory interventions relating to seatbelts, airbags, driver training, highway design, 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.

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 [61]; 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 [93]. It is a long-established principle in tort law that one should assign liability to the party best able to prevent the undesired outcome. Economics reinforces this: a party who can dump liability 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. It is also argued that a move to software liability could be harmful to the free software community. 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 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.

5.2 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 European country).

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 [87]. 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 interoperability 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 [99]. From a technical point of view, there is no reason to believe that secure and interoperable designs are mutually exclusive.

5.3 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 whom

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 [93]. 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* [42].

This 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.

Thus if a citizen buys a copy of Office in a shop, installs it on a PC, and suffers personal injury or property damage as a result, he can already sue. It is uncommon for such products to be used in safety-critical applications, although this is now starting; we have anecdotal evidence of spreadsheets being used to calculate dosages for radiotherapy and chemotherapy. The legal remedy available under the Directive is generally limited to injury and to personal property, in the sense of property intended and used mainly for private use or consumption. But this can be extended under other legal theories. Indeed, in the UK, the law on unfair contracts has been used to successfully bring a case, which we will describe below; that applied to a software error leading a local government to make an incorrect filing of tax data to central government. There are many other 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. The point we make here is that software vendors are not as immune from litigation as some advocacy groups paint them. A myth of 'immunity' may help dissuade some people from litigation, but it is a myth nonetheless. The real situation is more complex, and much more akin to the position with 'normal' goods such as cars and drugs.

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 citizen purchases a navigation system to use with a mobile home and, relying on it, is directed by a software error down a small country lane where his mobile home gets stuck, as a result of which he incurs significant towing and repair costs. 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 project⁴ 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 mobile home owner 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 in the consumer case. This suggests that, at least at the consumer level, 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.

The main issue with the Product Liability Directive is that it does not apply to business property. Thus although our mobile-home driver can sue, a truck driver whose lorryload of seafood got stuck and spoiled in exactly the same narrow lane has no recourse under the Product Liability Directive. Again, the Unfair Contract Terms Directive, or other legal doctrines, may come to the rescue. However, we should note that this complexity is a general problem for Community law rather than a specifically IT-policy one. There is a further problem in that a business might rely on software downloaded from a website in California, which makes it clear that the contract is governed by the laws of California, and subject to the exclusive jurisdiction of that state; if the contract contains an exclusion that is valid under California law then there may be little that the business can do if it is damaged by a software failure. Again, this is a general problem: in this case there may be little the Community can do, as even if EU courts took jurisdiction their judgments would not be enforceable in California.

⁴See <http://www.openstreetmap.org>

5.4 Software and systems liability options

This discussion should illustrate that software liability is both widely misunderstood and complex. But something may still need to be done. 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. Governments should not interfere in freedom to contract without good reason; and there is merit in the Microsoft point that software should not be singled out for unfair and discriminatory treatment, compared with hardware or services.

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 [93]. 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, for 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.

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.

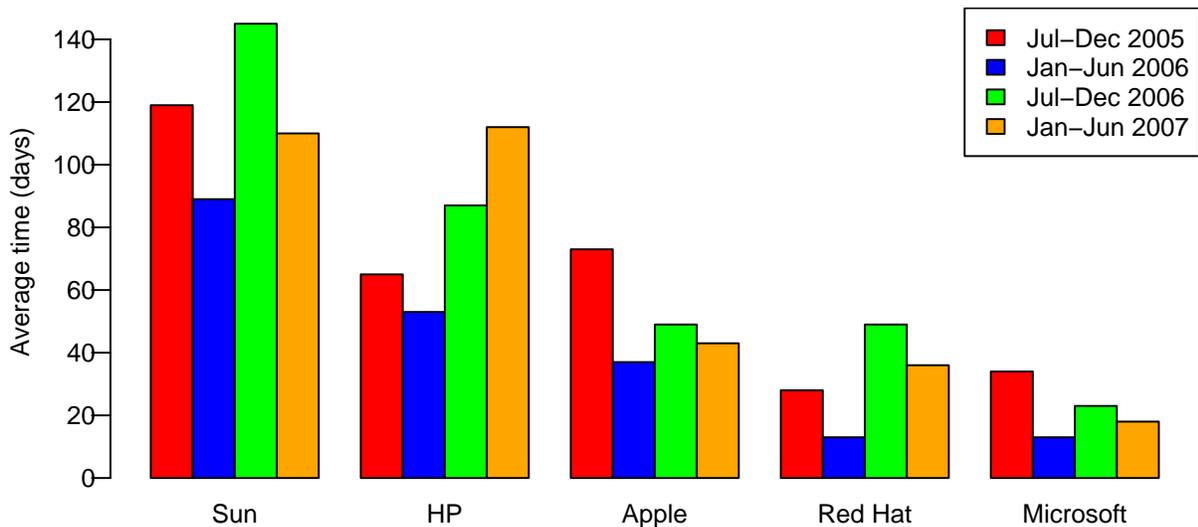
Recommendation 5: We recommend that the EU develop and enforce standards for network-connected equipment to be secure by default.

The precise nature of 'secure by default' will evolve over time. At present, the most important issue is whether the operating system is patched when the customer first gets it, and subsequently. One 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; the most likely would 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.

One of the stakeholders expressed concern at the likely costs if all consumer electronics required Common Criteria 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 on the facts. (We expect that once one or two landmark cases have been decided, the industry will rapidly adapt.)

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.



Source: Symantec ISTR vol. 9-12

Figure 12: Patch-development times for different operating systems

5.5 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. The publication of a patch often reveals the vulnerability to attackers, and 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.

5.5.1 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 5.5.1 indicates the time difference between the publication of vulnerabilities and the appearance of patches and exploits for vulnerabilities exploited by Chinese websites in 2007 [111]. 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

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. [111]

Table 1: 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.

wild before patches were available.

Nearly half of the vulnerabilities in Table 5.5.1 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 [97]. 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 [7]. 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 [18]. 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.

5.5.2 Challenge 2: Increasing patch uptake

While quantitative measurements are difficult to obtain, the view among security professionals is that patches are already available for the majority of exploits used by attackers. Over half of the exploits in Table 5.5.1 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 [8]. 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 [70]. 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 5.

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.

5.6 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. 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.

5.6.1 Fair contract terms

The main applicable law in the EU is the Unfair Contract Terms Directive [43], 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 [22]; 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 [102].

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 UK banks 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 [9]. The liability for fraud and security failure in online banking was thus transferred (at least on paper) to the customer. Yet there is significant variation across Member States in how complaints about fraudulent electronic banking transactions are handled. 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. By way of comparison, the US Regulation E, which governs electronic banking, places the onus of proof in practice 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 [48]). 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.

5.6.2 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’ [33]. 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 [31]. 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 [45]. As if that weren’t enough, spyware programs are often made deliberately hard to uninstall.

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, but there is evidence that the terms agreed between spyware vendors and US regulators are being flouted [34].

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 [32]. 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 [45] attempts to protect consumers from spam. For the most part, it prohibits sending any unsolicited messages to individuals, requiring their prior consent. However, Article 13 paragraph 5, 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.

5.6.3 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). These practices originated in the USA, which deregulated telecomms first, but are now creeping up the political agenda in Europe. The same goes for ‘identity theft’. Adam Shostack and Paul Syverson argue persuasively that identity theft is actually libel [100], and it’s spread from the USA to the UK. 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).

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. 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. 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 [47], 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, universal access to the Internet 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 [4]; 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.

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.

6 Dealing with the lack of diversity

Diversity can help security. Physical diversity deals with geographical distribution of redundant infrastructure components and network routes, whereas logical diversity means that distributed systems do not share common design or implementation flaws. A lack of diversity implies risk concentration which negatively affects insurability and thus an economy’s ability to deal with cyber risks.

6.1 Promoting logical diversity

For logical diversity to happen, alternatives must be widely available and adoption well-balanced. In practice, this has rarely occurred: technical lock-in, positive network externalities and high switching costs tend to yield dominant-firm markets [99]. 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. They are already on the agenda of the European Commission’s Interoperable Delivery of European eGovernment Services to Public Administrations, Businesses and Citizens (IDABC) initiative ⁵.

But even successful open standards do not always deliver diversity. Most applications supporting the Portable Network Graphics (PNG) format across platforms rely on the same reference implementation library *libpng* for image processing. As a result, vulnerabilities in one library (of which there are many: 17 vulnerabilities for *libpng*, including 5 critical, according to the National Vulnerability Database) can lead to multi-platform exploits. Hidden homogeneity at the lower levels can wreak havoc even when applications and systems platforms appear superficially diverse.

Option 2: Promoting diversity in procurement Consumers and firms are short-sighted when selecting a software product; the positive network externalities lead them to discount any increase in correlated risk. Governments need not be so myopic – but they often are. 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, strengthening the dominant platform [46]. And when citizens interact with their government online, they are often required to use Microsoft Office formats.

There have been several positive examples of governments choosing less dominant software platforms, albeit for cost-saving reasons. After a heated debate, the German Bundestag, the lower house of the federal parliament, decided in 2002 to move much of its server infrastructure to Linux and OpenLDAP [60]. 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. The French government spent 11 % of public IT expenditure on open source software in 2007 [101] and ran OpenOffice on 400,000 workstations [41].

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. However regulation tends to work rather more slowly than the industry. Cisco used to have a very dominant market position in the routers deployed in the Internet backbone. A vulnerability in Cisco routers [109] was disclosed that could remove a significant portion of the Internet backbone if a flash worm was disseminated. So the lack of diversity among routers used to be a critical concern. But 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

⁵see <http://europa.eu.int/idabc/>

Internet Explorer. In general, we feel the authorities should rather maintain a watching brief for competition issues that persist and have security implications.

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

6.2 Promoting physical diversity in CNI

Pitcom, a UK parliamentary group, has published a useful overview of CNI vulnerability aimed at legislators [91]. They show how an Internet failure could damage other parts of the CNI such as finance, food and health. Telecomms and power are known to be closely coupled: if a high voltage power line fails the engineers who go to fix it will keep in touch by mobile phone. But the mobile phones depend on the power supply to keep base stations operating. This particular problem can be fixed using satellite phones; what other problems should we anticipate?

6.2.1 Common mode failures and single points of failure

In principle, network designers avoid single points of failure using redundant components. However, as systems scale, they may be beyond an individual network's control even imagination. 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 administration systems for hospitals as far away as Cambridge and Great Yarmouth [103]. 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 [29], 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 [79].

6.2.2 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 there. An ISP with large numbers of customers will find it easy to arrange peering with an ISP with a large number of content providers because they can both avoid paying for transit. Thus companies such as Google, Akamai and the BBC will generally peer with anyone. But many ISPs are a hybrid of customers and content. In general, these hybrids can 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.

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. 11 EU countries have just one IXP; in almost all the others the largest IXP is 4 or more times the size 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. These pressures towards a dominant IXP lead to possible single points of failure at the IXP itself. Some leading IXPs have invested heavily in redundancy; others haven’t, mainly because of the expense.

For larger ISPs, an IXP failure 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. But 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, and it’s been counterproductive from the security viewpoint. The Access and Interconnection Directive [44] 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. It has made little or no difference to Internet transit provision, where a highly competitive market is keeping prices low. But it has reduced redundancy for Internet traffic flows. To avoid any risk of being caught by non-discrimination rules, large ISPs sometimes refuse to peer with small ISPs at foreign IXPs, in case they are forced by the regulator to provide free peering at their home IXP.

6.2.3 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 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 November 2005 the Commission adopted a green paper on a European Programme for Critical Infrastructure Protection (COM (2005) 576 final) [37].

The other, entirely European ‘hacking’ event of note is the April/May 2007 denial-of-service attack on Estonia. This event created widespread alarm, allegations 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 [74] 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 [52].

6.2.4 Policy options

Critical National Infrastructure is now understood to be a multi-national issue. One of the key difficulties in this area is that CNI companies do not wish to discuss how they might be vulnerable, while governments have limited understanding of the real world: for example the COCOMBINE project in Framework 6 examined IXPs but failed to understand why peering does or does not take place between particular ISPs, and merely attempted to find spatial patterns, with limited success [64, 57, 65].

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 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.

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.

7 Fragmentation of legislation and law enforcement

7.1 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.

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 [14] 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. But this solution does not look to be practical for cyberspace, because there is no global FBI.

The practical approach that has been taken is to try and harmonise national laws within a consistent international framework. The relevant treaty for the specific harms that cannot be dealt with by existing 'offline' legislation is the 2001 Convention on Cybercrime [24] 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, while 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. 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.

Other EU initiatives include a 2003 Council framework decision on attacks against information systems [25] which required the offences of 'illegal access to information systems', 'illegal system interference', 'illegal data interference' along with 'instigation, aiding, abetting and attempt' to be criminalised by 2005; and the 2007 Commission draft communication on cybercrime [40], which defined cybercrime as traditional crimes committed over electronic networks, illegal content (child abuse pictures, etc), and 'crimes unique to electronic networks'. (Its section on legislation was vague.)

7.2 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.

7.2.1 Defining the problem

Police forces must make tough choices in deciding which crimes to investigate. In the case of electronic crime, one of the first questions is how many local citizens are affected, and how many local computers are being used to launch attacks. Using this criteria, most attackers are not worth pursuing, even if in aggregate they are having a devastating

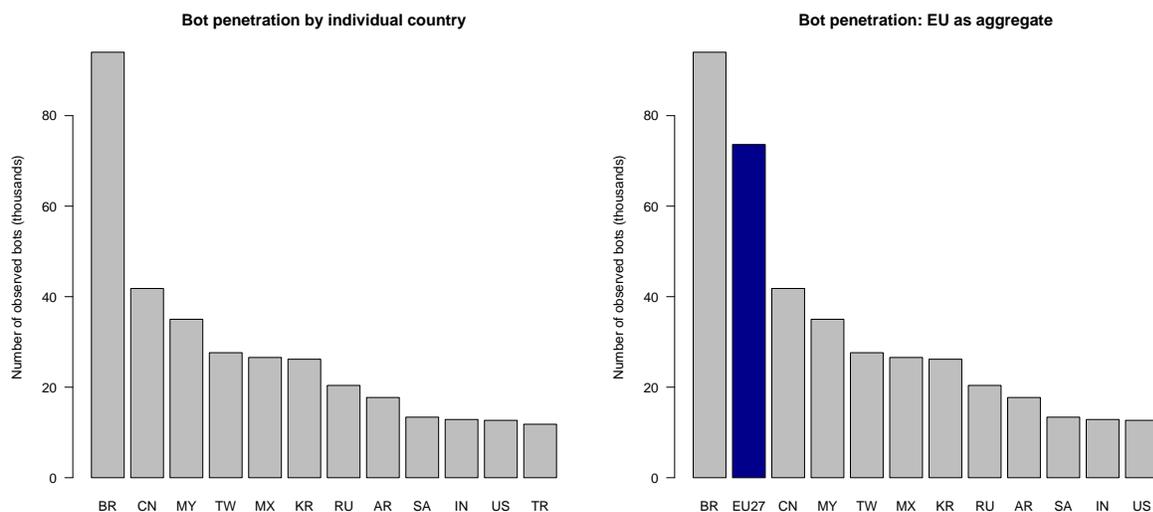


Figure 13: Not our problem? Number of global botnet victims identified by the Chinese Honeynet Project between June 2006 and December 2007. No European country is among the top dozen alone (left), but the European Union as a whole is second only to Brazil (right). *Source:* Own aggregation based on data of [110]

effect (see Figure 13). Even those cases that are deemed worth pursuing almost always lead to computers located in other countries. The current structures for international co-operation were designed for physical crimes, where cross-border activity is rare. They slow down investigations and drive up costs. As a result, very few cyber-criminals are successfully prosecuted. This makes attacks more attractive and therefore more prevalent.

7.2.2 Methods for co-operation

There are several two traditional options for law enforcement agencies when a crime involves 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, 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 force. A major difficulty is that it is hard to predict the cost prior to approving the operation. Joint operations are largely unfunded and carried out on a quid pro quo basis, so they cannot be relied upon as the baseline 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 Co-operation may also 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,

and are very slow to process. So many investigators prefer to avoid using them where possible.

Option 3: Cyber-security co-operation using NATO as a model 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. Part of this budget is then used to fund liaison officers at a central command centre. That command centre decides what tasks are to be undertaken – and the liaison officers relay requests to their own countries’ forces. This is in effect a permanent ‘joint operation’, and avoids the glacial speed of MLAT arrangements. The key is that countries trust their liaison officers to assess which requests carry no political baggage and can be expedited at once.

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.

8 Other issues

8.1 Cyber-insurance

Cyber-insurance is a common tool for cyber-risk management, in particular to transfer residual risk [59, 10, 75]. A brisk market is 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. However, the practice looks a bit different. Firms often find their premiums based on non-technical criteria such as firm size or individual loss history.
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. However, we are aware of only one concrete case in which an insurance association funded original research on vulnerabilities.
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. This reduces their reserve requirements.
4. **Market-based security metric.** As discussed earlier, insurance premiums may serve as market-driven metrics to quantify security [59]. 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.

Yet the market appears to perform below expectations. The USD 350 million estimated global market size in 2005 [26] is only one-tenth of a forecast made for 2005/06 by the Insurance Information Institute in 2002 [71] and below one fifth of a revised forecast from 2003 [67]. 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 years⁶ 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.

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 [10]. Other 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 [26]. 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 help establish a wider market for cyber-insurance might conceivably be justified. Several options are possible in theory.

⁶28% in 2005, 25% in 2006 [23]

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. Finally, there has been a strong lobbying effort from the anti-virus industry to make their own products a compulsory purchase; if governments are going to compel, why not take this more direct route?

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 [30]. 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. 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 [76] 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 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. We mentioned Exploit derivatives; Cat bonds [28] are another class of instruments for insurance risk securitization, whose pay-out function is defined on actual impact rather than on the market's assessment of the future probability of a breach. There is some experience with cat bonds in flood and natural-disaster insurance. A difficulty in applying them 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-outs 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 [73, 88, 20, 10, 12, 13]. 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’ [13]

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 [10]. 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 We were unable to make 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.)

8.2 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. Many small software developers are unaware of this control regime and may be technically in breach of its implementation provisions in some Member States. More recently, in some Member States, well-meant but poorly drafted legislation has impeded security research. 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 [3]. 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 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 [21].

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.

9 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.

In our report, of which this is an abridged version, we 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 with the authors alone.

References

- [1] 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)
- [2] Akerlof, G. A. (1970): The market for 'lemons': quality uncertainty and the market mechanism. *Quart. J. Economics*, **84**, 488–500
- [3] 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>
- [4] Anderson, R. J. (2001): *Security Engineering – A guide to building dependable distributed systems*, Wiley. <http://www.cl.cam.ac.uk/~rja14/book.html>
- [5] Anti-Phishing Working Group: <http://www.antiphishing.org/>
- [6] APACS (2007): Card fraud losses continue to fall. Press Release, 14 March. http://www.apacs.org.uk/media_centre/press/07_14_03.html
- [7] 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.
- [8] 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
- [9] 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/

- [10] Böhme, R. (2005): Cyber-insurance revisited. *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA. <http://infosec.net/workshop/pdf/15.pdf> (last access: 13 Nov 2007)
- [11] 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
- [12] 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)
- [13] 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)
- [14] Brenner, S (2007): Bonnie & Clyde and cybercrime. <http://cyb3rcrim3.blogspot.com/2007/11/bonnie-clyde-and-cybercrime.html> (last access: 10 Jan 2008)
- [15] 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
- [16] 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
- [17] Casper, C. (2007): Examining the feasibility of a data collection framework. *ENISA Technical Report*.
- [18] 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)
- [19] Cavusoğlu, H., Mishra B. and Ragunathan, 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
- [20] 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.infosec.net/workshop/pdf/47.pdf> (last access: 3 Dec 2007)
- [21] 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/>
- [22] Collins, B. St. J. (1995): Unfair terms in consumer contracts regulations 1994, at <http://webjcli.ncl.ac.uk/articles3/collins3.html>
- [23] Computer Security Institute (2007): *The 12th Annual Computer Crime and Security Survey*. <http://www.gocsi.com/> (last access: 8 Nov 2007)
- [24] 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)

- [25] 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)
- [26] 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)
- [27] 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
- [28] D'Arcy, S. and France, V. G. (1992): Catastrophe futures: A better hedge for insurers. *J. Risk and Insurance*, 59 (4), 575–600
- [29] 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)
- [30] Duffy, D. (2002): Safety at a premium. *CSO Magazine*, December. <http://www.csoonline.com/read/120902/safety.html> (last access: 3 Dec 2007)
- [31] Edelman, B. (2004): 180solutions installation methods and license agreement. <http://www.benedelman.org/spyware/180-affiliates/installation.html> (last access: 18 Dec 2007)
- [32] Edelman, B. (2007): Advertisers using WhenU. <http://www.benedelman.org/spyware/whenu-advertisers/> (last access: 18 Dec 2007)
- [33] Edelman, B. (2007): Spyware: Research, testing, legislation, and suits. <http://www.benedelman.org/spyware/> (last access: 18 Dec 2007)
- [34] 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)
- [35] van Eeten, M. J. G. et al. (2007): *The Economics of Malware: Security Decisions, Incentives and Externalities*. Draft OECD report.
- [36] 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)
- [37] 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)
- [38] 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)
- [39] 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)

- [40] 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>
- [41] European Communities (2006): French administration opts for OpenOffice. 7 July 2006. <http://ec.europa.eu/idabc/en/document/5695/469>.
- [42] 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)
- [43] 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
- [44] 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>
- [45] 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>
- [46] 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>
- [47] 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>
- [48] 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>
- [49] 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
- [50] Fama, E. (1970): Efficient capital markets: A review of theory and empirical work. *J. Finance*, **25** (2), 383–417
- [51] 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)
- [52] Fox News (2008): Estonia Charges Solo Hacker for Crippling Cyberattacks. Jan 25 2008. <http://www.foxnews.com/story/0,2933,325547,00.html>

- [53] 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
- [54] 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)
- [55] Gal-Or, E. and Ghose, A. (2005): The economic incentives for sharing security information. *Information Systems Research*, **16** (2), 186–208
- [56] Garg, A., Curtis, J. and Halper, H. (2003): Quantifying the financial impact of IT security breaches. *Information Management & Computer Security*, **11** (2), 74–83
- [57] Giovannetti, E., Neuhoff, K. and Spagnolo, G. (2005): Agglomeration in Internet cooperation peering agreements. *Cambridge Working Papers in Economics*, 0505. <http://econpapers.repec.org/paper/camcamdae/0505.htm> (last access: 10 Jan 2008)
- [58] 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
- [59] Gordon, L. A., Loeb, M., Sohail, T. (2003): A framework for using insurance for cyber-risk management. *Comm. ACM*, **46** (3), 81–85
- [60] Heise Online (2002): Tux takes its seat in Germany’s federal parliament. 28 February, at <http://www.heise.de/english/newsticker/news/25255>
- [61] 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)
- [62] 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
- [63] 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
- [64] 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>
- [65] 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)
- [66] 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)
- [67] Insurance Information Institute (2003): Computer security-related insurance issues. <http://www.iii.org/media/hottopics/insurance/computer/> (historical state: 01 Oct 2004)
- [68] Kannan, K. and Telang, R. (2005): Market for software vulnerabilities? Think Again. *Management Science*, **51** (5), 726–740

- [69] Keizer, G. (2007): Gartner: Bug bounty hunting is 'risky endeavour', *Computerworld* (2 May), at <http://www.computerworlduk.com/management/security/cybercrime/news/index.cfm?newsid=2836>
- [70] 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>
- [71] 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.infosecnet.net/workshop/pdf/42.pdf> (last access: 3 Dec 2007)
- [72] Knuth, D. (2002): All questions answered. *Notices of the AMS*, **49** (3), 318–324 <http://www.ams.org/notices/200203/fea-knuth.pdf>
- [73] Kunreuther, H., Heal, G. (2003): Interdependent security. *Journal of Risk and Uncertainty*, **26** (2/3), 231–249
- [74] Lesk, M. (2007): The new front line: Estonia under cyberassault. *IEEE Security and Privacy*, **5** (4), 76–79
- [75] 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)
- [76] 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)
- [77] McAfee Inc. (2007): *McAfee Virtual Criminology Report*. http://www.mcafee.com/us/research/criminology_report/default.html (last access: 8 Dec 2007)
- [78] McPherson, D., Labovitz, C., Hollyman, M. (2007): *Worldwide Infrastructure Security Report*, vol. 3, Arbor Networks. <http://www.arbornetworks.com/report>
- [79] 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
- [80] 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)
- [81] 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)
- [82] 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
- [83] Mozilla Corporation (2007): Mozilla security bug bounty program. <http://www.mozilla.org/security/bug-bounty.html> (last access: 22 December 2007)

- [84] 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)
- [85] Nader, R. (1965): *Unsafe at Any Speed*. Grossman Publishers Inc., New York.
- [86] National Conference of State Legislatures (2007): Breach of information. <http://www.ncsl.org/programs/lis/cip/priv/breach.htm>
- [87] Office of Fair Trading (2003): Payment systems. http://www.oft.gov.uk/advice_and_resources/resource_base/market-studies/payment-systems
- [88] Ogut, H., Menon N., Ragunathan, 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.infoecon.net/workshop/pdf/56.pdf> (last access: 3 Dec 2007)
- [89] OpenDNS (2007): OpenDNS shares April 2007 PhishTank statistics, Press Release, 1 May. http://www.opendns.com/about/press_release.php?id=14
- [90] 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)
- [91] 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)
- [92] Privacy Rights Clearinghouse (2005): A chronology of data breaches. <http://www.privacyrights.org/ar/ChronDataBreaches.htm>
- [93] Rice, D. (2007): *Geekonomics – The Real Cost of Insecure Software*. Addison-Wesley, New York.
- [94] Sans Institute (2007): SANS Top-20 2007 security risks. <http://www.sans.org/top20/>
- [95] Schechter, S.E. (2004): *Computer Security Strength & Risk: A Quantitative Approach*. PhD thesis, Harvard University, Cambridge, MA
- [96] Schwalb, M. (2007): Exploit derivatives & national security. *Yale J. Law and Technology*, **9**, 162–192
- [97] Security Focus Inc. (2007): Bugtraq mailing list. <http://www.securityfocus.com/archive/1> (last access: 6 Dec 2007)
- [98] 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)
- [99] Shapiro, C., Varian, H. R. (1999): *Information Rules. A Strategic Guide to the Network Economy*. Harvard Business School Press, Boston, MA
- [100] 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

- [101] 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-1-open-source-va-prendre>
- [102] St Albans District Council vs ICL (1996), at http://www.smaldonado.com/marcos/docs/it_case_su_uk_en.html
- [103] 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)
- [104] 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)
- [105] Symantec (2007): *Internet Security Threat Report*. <http://www.symantec.com/business/theme.jsp?themeid=threatreport> (last access: 6 Dec 2007)
- [106] 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)
- [107] 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)
- [108] Wolfers, J. and Zitzewitz, E. (2004): Prediction markets. *J. Economic Perspectives*, **18** (2), 107–126
- [109] 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)
- [110] 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)
- [111] Zhuge, J., Holz, T., Song, C., Guo, J., Han, X., Zou, W. (2008): Studying malicious websites and the underground economy on the Chinese web. *Workshop on the Economics of Information Security (WEIS)*, Dartmouth College, Hanover, NH.