



Quarterly

12/2005

IN THIS EDITION

	Page
A Word from the Executive Director	1
A Word from the Editor	2
From the World of Security – A Word from the Experts	3
Patching Strategies	3
Trends in Security Economics	6
Security and the European Standards Organisations	8
Past ENISA events	10
NIS: Political and Technical Challenges – Rome	10
Readiness for Handling NIS Incidents – Lithuania	10
Secure 2005 – Poland	12
BSI Information Security Management Workshop – Bonn	14
National Liaison Officers Day	14
From the Member States	16
Hungary's National NIS Projects	16
SEMA Visits ENISA	18
BSI Materials for IT Security in Companies in Critical Infrastructures	19
The EBIOS Method	19

A WORD FROM THE EXECUTIVE DIRECTOR



Dear Readers,

As we close 2005, it is useful to reflect on our first year as an Agency and on the state of information security in general.

Time certainly does fly. It has been five years since the heyday of the dotcom boom. Despite the lull that followed, the Internet has never stopped evolving and has become increasingly interwoven with the fabric of our daily lives. From shopping online to instant messaging with friends, more and more of our daily activities involve the Internet.

Indeed, while most people reading these pages can probably still recall the first time they sent an e-mail or visited a webpage, for many of the teenagers of today this is a completely natural part of life; for them life without the Internet would be unthinkable.

At the same time the unabated growth of the Internet, online communities, and online commerce has also led to a myriad of new threats and security risks. And while there is certainly no need for panic or undue fear, it is clear that added vigilance and measures are needed in the face of these risks.

How can users stay secure and avoid the pitfalls of phishing attacks, identity theft, and other security incidents? Many of the rules for online safety are the same as those offline; for example, to avoid many Internet scams, users would be well advised to remember the old adages that there is no free lunch, and that if something sounds too good to be true, then it probably is.

And yet user vigilance is not enough - unfortunately even the most cautious of users can fall victim to crimes that exploit technical or procedural vulnerabilities in increasingly complex systems. This necessitates a collective effort to ensure that security does not become the spoiler in our move towards a digital society. This must involve all players including the private sector, users, and public organisations.

ENISA was set up with the aim of addressing these challenges within its well defined scope, and in particular to provide a uniquely European approach that focuses on information exchange, co-operation, and learning from each other.



Trends in Security Economics

Tyler Moore, Ross Anderson



Tyler Moore



Ross Anderson

One of the most exciting and rapidly growing fields of research today is the economics of information security. Many security and privacy failures are not purely technical; rather, misaligned incentives are often to blame. For instance, people in the best position to protect a system may be poorly motivated if they do not have to bear the costs of failure. This article explores two of the key economic challenges to information security – incentives and metrics. We then discuss one of the applications, namely assessing the merits of open versus closed systems, and finally we highlight several promising research areas.

Misaligned incentives

It is an established principle of civil law that liability should be assigned to the entity best placed to manage risk; otherwise, perverse incentives may preclude the adoption of reasonable countermeasures. This has often been ignored by information security system designers. A classic example is ATM fraud. When a customer disputes a transaction, banks in the United States must demonstrate the customer is mistaken or lying to avoid responsibility. As a result, US banks have long been motivated to defend their systems properly. However, in Britain, Norway and the Netherlands, the burden of proof used to lie with the customer – creating a really hard problem for a victim of fraud. Banks underinvested in protection mechanisms until higher rates of fraud became untenable, and the liability was reassigned. Similar examples abound. Distributed denial of service attacks exploit unprotected home and university computers to target popular websites. So long as the attack does not pose a direct threat to the end user, few security measures will be adopted.

However, dumping liability on end users achieves little, since most users do not have the expertise to defend against such attacks. Many ISPs and mail service providers offer spam filtering services. More recently, liability has started to shift in practice to network operators. Their main incentive is that networks which produce large volumes of spam place their peering arrangements at risk.

Many open questions remain. We do not yet know the optimal balance between technical and regulatory mechanisms. How much reliance should we place on prevention, and how much on filtering? Should vendors of insecure platforms be liable? How much filtering should be done on egress from source networks, how much by the destination networks, and how much by the destination mail systems? What is clear is that any viable solution will have to consider both technology and incentives.

Measuring security strength: markets, auctions and return-on-investment

One common factor in many computer security problems is the existence of informational asymmetries. Akerlof won a Nobel Prize in Economics for explaining the pitfalls of a market with asymmetric information: if buyers cannot readily determine the quality of goods on offer, then they will not pay a premium for a high quality version. This drives out producers of quality goods, since they will not be compensated for their efforts. Used cars are a classic example – it is often impossible for a buyer to tell the difference between a good used car and a “lemon”. The market price then falls to the price of unreliable cars, which undermines sales of good cars.

The software market too is a market for lemons. Because there is no good way for consumers to tell secure products from insecure ones, companies do not invest in product security. Certification schemes such as the Common Criteria were an early response; more recently, researchers are looking for ways to use markets to elicit information. Early attempts included vulnerability markets, where the current market price for an undiscovered exploit indicates the current level of system security. Refinements include establishing vulnerability auctions and collecting historical data to help predict the rate of subsequent discoveries.

Open versus closed systems

One of the hottest debates in security has centred on open versus closed systems. Are open systems such as GNU/Linux more secure than proprietary offerings such as Windows?

Here, too, security economics can provide some valuable guidance. If security vulnerabilities are uncorrelated with each other, then open and closed systems would be equivalent: opening up a system would make bugs easier for both the attacker and the defender to find and would thus help them both equally. The problem therefore becomes an experimental one: are security vulnerabilities correlated or not? Early statistical evidence suggests that they are. A recently discovered vulnerability is more likely to be rediscovered by someone else than one would expect from random chance.

Security economics can also throw light on the behaviour of markets that are open, or closed. For example, Hal Varian presented a surprising result at a Digital Rights Management (DRM) conference in January 2005 – that stronger DRM would help platforms more than the music industry, because the computer industry is more concentrated (with only three serious DRM suppliers – Microsoft, Sony and the dominant firm, Apple). Before the end of the year, the music industry was protesting that Apple was getting an unreasonably large share of the extra revenue being generated by online music sales. This surprised industry observers; the music industry in particular had expected that stronger DRM would help it make more money. It was a striking demonstration of the predictive power of economic analysis when applied to technical security mechanisms.

Towards a network perspective

Network analysis can be a powerful tool for information security economists. Networks are everywhere, from academic citation



patterns to terrorist cell communication structures to underground file-sharing darknets. Each of these networks has distinct structural properties – from the density of connections between neighbours to the centrality of particular nodes. In fact, while much progress has been made in identifying the distinguishing characteristics of networks, the relationship between these properties and resulting attack and defence strategies is not at all well understood.

Social scientists use network analysis to examine the complicated interactions of large, decentralised groups in human society. But these tools have striking relevance for computer scientists as well. Decentralisation makes possible computer networks without unified control. The BGP (Border Gateway Protocol) routing that underlies the Internet is one example; peer-to-peer networks are a second; and the ad-hoc networking techniques under development by the sensor-network

these questions and discuss the direction solutions to others might take.

Several peer-to-peer systems, from the Eternity Service to Chord, distribute content randomly throughout the network. Yet for other systems, including most file-sharing applications like Gnutella and Kazaa, users are allowed to choose which content to share. Offering nodes discretion motivates them to spend more on defence. This is because individual preferences do not necessarily align with those of the larger community. A user with left-leaning political views may quickly lose interest in protecting shared information if he must serve right-wing content too. Yet a successful defence depends on the collective efforts of the system's users. Providing choice can maximise individual utilities as well as individual incentives to invest in defence. Future large-scale systems may be more a federation of networks, or of clubs, than large homogeneous entities.

Much has been made of the positive externalities generated by interconnection in networks. For instance, Metcalfe's Law claims that the value of a network is proportional to the square of its size. Yet there are also costs associated with network growth, notably security-related ones. In fact, without a single entity in control, the marginal costs of adding new nodes can actually rise with the network size. Peer-to-peer networks often grow to the point that free-riding cannot be mitigated effectively: the benefits brought by new nodes (e.g., files for sharing) no longer outweigh the security costs of new members. Over-participation then destroys the network. (An early example was CB radio, whose utility was undermined by congestion.) An interesting challenge is to capture these costs and examine the resulting impact on network formation.

We suspect that the scalability of security is a fundamental limiting factor in network growth, and indeed in the ways in which technology and society interact. A village is different from a city, and a global network is something else again. When designing dependable infrastructures for networked global society, we need to understand these issues better.

Conclusion

The economics of information security started out when we recognised that systems often fail for non-technical reasons. Security economics remains a vibrant cross-disciplinary effort, bringing in people from a wide variety of backgrounds, and throwing up deep and fascinating new challenges. Starting from the analysis of operational questions such as the right level of security investment and the incentives facing security managers, it has developed through the analysis of controversial questions such as the merits of open versus closed systems to such basic issues as the ways in which networks – and the institutions they support – scale in the face of accident, error and attack.

Until now, the field's annual event – the Workshop on the Economics of Information Security (WEIS) – has been held in the USA, but in 2006 it will come to Europe for the first time. To learn more about security economics, mark your diary for WEIS 2006 in Cambridge, England, on 26-28 June.

Tyler Moore is working for a PhD in security economics at the Computer Laboratory, Cambridge University, England.

Ross Anderson is Professor of Security Engineering at Cambridge University.



research community provide a third. Nodes can choose their behaviour, their transaction partners and even whether to participate in the system. The difficulties imposed by asymmetric information extend naturally to such networks, since individual nodes cannot view a network's complete topology or observe all its communications.

Many questions pertinent to security appear. Is it better for defenders to aggregate or disperse in the face of censorship attacks? Which network topologies are most resilient to targeted attack? What socially optimal network structures might balance the costs of maintaining security against the advantages of interconnection? And how do networks formed in practice differ from the best case? We can offer answers to some of

For many networks, a small number of key players are critical to operational success. As a result, adversaries seeking to undermine the network target these nodes for attack. Music industry agents attempting to disrupt peer-to-peer file-sharing networks target individuals believed to have been running major nodes, using techniques from technical attacks to aggressive litigation. The same principle works in politics: as a handful of leading individuals often do much of the work to hold a society together, subverting or killing these leaders is likely to be the cheapest way to make an invaded country submit. But how might the defenders react? What are the optimal strategies of attack and defence? And how do defence tactics change under imperfect information, where co-ordination is imprecise?