Ethics Inf Technol (2014) 16:227–240
DOI 10.1007/s10676-014-9348-3

ORIGINAL PAPER

A roadmap towards improving managed security services from a privacy perspective

Nils Ulltveit-Moe

Published online: 5 August 2014

Abstract This paper proposes a roadmap for how privacy leakages from outsourced managed security services using intrusion detection systems can be controlled. The paper first analyses the risk of leaking private or confidential information from signature-based intrusion detection systems. It then discusses how the situation can be improved by developing adequate privacy enforcement methods and privacy leakage metrics in order to control and reduce the leakage of private and confidential information over time. Such metrics should allow for quantifying how much information that is leaking, where these information leakages are, as well as showing what these leakages mean. This includes adding enforcement mechanisms ensuring that operation on sensitive information is transparent and auditable. The data controller or external quality assurance organisations can then verify or certify that the security operation operates in a privacy friendly manner. The roadmap furthermore outlines how privacy-enhanced intrusion detection systems should be implemented by initially providing privacy-enhanced alarm handling and then gradually extending support for privacy enhancing operation to other areas like digital forensics, exchange of threat information and big data analytics based attack detection.

Keywords Security · Privacy · Outsourcing · Intrusion detection and prevention systems · Managed security services · Ethical awareness

N. Ulltveit-Moe (✉)
University of Agder, Jon Lilletuns vei 9, 4879 Grimstad, Norway
e-mail: nils.ulltveit-moe@uia.no

Introduction

Various attack detection techniques, like Intrusion Detection Systems (IDS), spam filters or anti-virus are being used to detect, investigate or prevent cyber-crime both in the private and public sector. It is legal to perform monitoring of computer networks and hosts using such potentially privacy invasive technologies in most European countries, as long as the *purpose* with the monitoring is to detect cyber-attacks. There are for example explicit exceptions for measures related to detecting cyber-attacks in the EC communications directive (European Commission 2002). Such monitoring may nevertheless be problematic from a privacy or confidentiality perspective, because the *effect* of such monitoring is largely unknown. This means that better methods are needed to protect private or confidential information, at the same time as better techniques are required for ensuring transparency on use of such information.

This paper discusses privacy and confidentiality problems related to Managed Security Services (MSS) which are security monitoring services that have been outsourced to a service provider (Kairab 2005). The main focus in this paper is on attack detection techniques like Intrusion Prevention Systems (IPS) and Intrusion Detection Systems (IDS). Based on this ethical discourse, a roadmap is proposed on how to improve handling of private and confidential information for such systems.

The paper is organised as follows: The next section gives an introduction to what intrusion detection and prevention systems are. It also discusses the effect of outsourcing security monitoring, as well as giving a definition of privacy. Section 3 investigates how privacy can be improved for IDS from a high-level perspective. Section 4 covers sources of information leakages in intrusion detection and prevention

systems. It also discusses the conflict of interest between privacy and security from an economic perspective, where both compete for the same funding. Section 5 discusses privacy valuation and Sect. 6 investigates advantages and disadvantages with automatic attack prevention (IPS) compared to attack detection (IDS) from an ethical perspective. Section 7 proposes a roadmap towards improved privacy for managed security services based on ethical principles, technical privacy enforcement mechanisms, privacy metrics and best practices within information security management. Section 8 concludes the paper and outlines future work.

What is intrusion detection and prevention?

Network based Intrusion Detection System (IDS) is the Internet equivalent of a burglar alarm which monitors all packets passing through a router, switch or firewall. The two main types of IDS are signature-based IDS, which relies on matching known attack patterns in the data traffic, and anomaly-based IDS, which interprets statistical anomalies in the data traffic as possible attack activities. Intrusion Prevention Systems (IPS) extend IDS with functionality to automatically block traffic from senders that triggers an IDS rule or traffic anomaly. Signature-based IDS is typically performed using *deep packet inspection* (DPI), which means that the following data can be investigated: packet header information, e.g. IP addresses and ports; payload in each data packet; reassembled streams of data spanning several data packets; and entire communication sessions between a client machine and a server.

This means that all communication between a client application and a server in principle can be intercepted and investigated in detail by an IDS, as long as the communication sessions are not encrypted. However encrypted protocols can also be monitored using application level intrusion detection systems. One such example is ModSecurity which works on Apache, IIS7 and Nginx web servers (Trustwave 2014). Modern intrusion detection systems can also monitor encrypted TLS/SSL traffic, given that the IDS is trusted with the digital certificate of the connection endpoint (Plashchynski 2014). This means that it is technically possible for an organisation to monitor any information that goes through the company's own services. It is however not in general possible to monitor encrypted traffic towards third party organisations, since the intrusion detection system then would lack the necessary keys to decrypt the information.

The fact that *any* information from a provided service in principle can be monitored, if the owner of a service has a desire to do so, means that extra care is needed to avoid leaking private or confidential information if such monitoring is being outsourced to third-party organisations. This is in particular important for information that is inherently sensitive, for example Personally Identifiable Information and patient information in hospital information systems, personal data in traffic control systems or confidential data in critical infrastructures. Private or confidential information must for such systems be protected by encryption from inception and until the data is safely destructed according to the Privacy by Design principles (Cavoukian et al. 2010). This means that nobody without authorisation, not even system administrators, should be able to access this information.

The market leaders/challengers for IDS/IPS devices: Cisco, HP, IBM, Juniper, McAfee and Sourcefire (Gartner 2010), all rely on using DPI, based on a combination of signature-based and anomaly-based detection techniques. This avoids any security blind-spots that may occur using either technology (Cisco 2013; Roesch and Green 2009; Bicknell and Jean 2011; X-Force 2011; McAfee 2007). IDS focuses on identifying possible incidents, and supports incident response efforts to identify successful compromise of a system due to an adversary exploiting a system vulnerability (Scarfone and Mell 2007). Typical use of IPS technology involves acting as a shield protecting vulnerable machines from known attacks (Gartner 2010). This is in particular important for Critical Infrastructures, which may have long patch latencies due to strict safety requirements on testing of patches before deployment. Other uses of IDS include identifying security policy problems, documenting the existing threat to an organisation and deterring individuals from violating security policies (Scarfone and Mell 2007).

The privacy concern related to IDS or IPS comes both from policy-based IDS/IPS rules and from false alarms or anomaly patterns in attack detecting rules which may leak Personally Identifiable Information (PII) or other confidential information. These privacy leakages may also occur due to activities caused by malicious actors, for example computers compromised with malicious software that reveals sensitive information about compromised users. The latter may cause privacy leakages both towards the malicious actors, which should be detected and deterred by IDS/IPS, as well as to MSS providers monitoring the IDS alarms.

The Open Source intrusion detection system Snort is a technology where the amount of rules performing deep packet inspection can be quantified. Snort is one of the market leaders within intrusion detection and prevention systems (Gartner 2010). Investigating the commercial VRT rule set for Snort shows that 99.7 % of the IDS rules that by default are enabled (4,503 out of 4,517 rules at time of writing) contain the "content:" directive, meaning that

DPI is being heavily used for investigating the content of packets for Snort (Sourcefire Vulnerability Research Team VRT 2014). Signature-based IDS will in general depend on using deep packet inspection for detecting attack vectors, since the IDS signatures need to look for data patterns that resemble known attack vectors in the payload data.

One could perhaps still argue that the Snort numbers would not generalise to other IDSs, however this goes beyond the main point: Signature-based intrusion detection systems both can and will use DPI for identifying attacks. There is therefore currently a *big uncertainty* on what *impact* such surveillance has on privacy and confidentiality, since there is no transparency on what is being measured as well as no convincing way to audit whether the surveillance is strictly necessary or not.

It should therefore be the computer security industry that needs to convince the public that such surveillance indeed is needed, and not vice versa. The public should have a right to know what such surveillance entails, as well as trustworthy evidence that such monitoring is being performed according to a strict definition of need. The public should also be assured that the data from such services are being adequately protected, so that data is not being kept longer than necessary and that only the necessary amount of people can investigate possible security incidents. This means that secure auditable logging routines as well as privacy metrics are needed for such operations.

The effect of outsourcing security monitoring

Outsourcing security monitoring to MSS providers has gained popularity for two main reasons. First, the cost of providing 24×7 monitoring is only a fraction of what such monitoring would cost in-house (Ding et al. 2005). Second, MSS providers have in general got more experience in handling security incidents and more updated monitoring technology, by specialising in this area, than the average customer. A large client base also contributes to service quality improvements because an MSS provider, monitoring a large set of networks, easier can correlate attacks and identify new attack patterns. They can also share information about attacks and attack mitigation strategies between its customers, which is one of the factors that have been shown to reduce the risk of attacks from adversaries (Schechter and Smith 2003). One concern firms have when considering to outsource security services, is that the MSS provider may shirk (avoid doing its duties) secretly to increase profits. In economics this behaviour is commonly referred to as the Moral Hazard problem. The optimal way to avoid such behaviour on a contractual basis is to use a performance-based contract, however the degree of performance dependence may decrease if the reputation effect becomes significant (Ding et al. 2005).

It should be noted that the Moral Hazard problem not only is applicable to the security of the monitored data. It is also applicable to the privacy and confidentiality of the monitored data. Both in the sense of handling more private and confidential information than strictly necessary and in the sense of potentially leaking or abusing private or confidential information. This does in the end mean that a principal (here the customer of security services) should require that both the *security* and *privacy* performance for outsourced MSS should be part of a performance-based contract with the MSS provider. This means that the MSS provider should be *accountable* for both the privacy and security part of the operation which means that suitable performance metrics and activity logging procedures are needed for both privacy and security, so that the performance in these areas can be reported and audited if necessary. This is in line with the 6. foundational Privacy by Design principle, that the privacy-enhanced design must ensure transparency (Cavoukian et al. 2010).

What is privacy?

Privacy is a broad concept that can have different meaning in different contexts. Warren and Brandeis early on defined privacy from a legal perspective as *the right to be let alone* (Warren and Brandeis 1890). Other definitions focus more on privacy as an intellectual property from a utility perspective, where the data owner should be ensured *self determinism* about private data (Samuelson 2000). This amongst others means that the data owner must be able to give and revoke consent to access private data (Cavoukian et al. 2010). This has for example lead to actuarial models that aim at estimating the perceived cost of privacy leakage for insurance contracts (Gritzalis et al. 2007).

This paper considers privacy or confidentiality as an intellectual property that has a subjective value by the information owner, and therefore should be protected from unnecessary disclosure. This furthermore means that access to such information should be transparent and auditable.

How can privacy be improved?

One way to illustrate how technology can improve privacy from a high level perspective is airport security. Many are willing to trade some convenience and privacy for added security. It is therefore accepted in our society that all passengers undergo privacy-invasive security control checks when travelling by airplanes to increase the perceived safety. The privacy-invasive security controls aim at reducing the possibility that adversaries, like terrorists or psychologically unstable persons, bring weapons, explosives or other dangerous items on board the airplane.

There has been quite extensive research on more efficient ways to detect hidden weapons on people. One efficient technology, that recently has been deployed, is backscatter X-ray scanners (Cavoukian 2009). These scanners expose the person to be checked with small amounts of X-ray radiation, and use the backscatter X-rays to produce photo-quality images that can see through clothes. This technology is used as an alternative to personal searches, since it easily can reveal hidden weapons. If a suspicious item is detected, then the security officer will perform a manual search to verify what the suspicious item is.

This technology causes a privacy concern, since it essentially shows a naked picture of the person being scanned. Privacy enhancing technologies have therefore been implemented to deal with this problem. The techniques include using blurred pictures or stylistic images, emphasising items that are not considered normal body features. Such techniques mean that the privacy of all people who are not being suspected of carrying illegal items need not be violated, which limits the amount and degree of privacy violations.

The Internet analogy of this is surveillance techniques like IDS using deep packet inspection. This means that the MSS provider effectively can see any cleartext traffic that triggers IDS alarms between a customer performing a service on the Internet and the service provider. There may therefore be a conflict between the privacy and security objectives[1] for managed security services. However, a larger problem is the lack of transparency on what is being monitored, and why. My experience is that IDS rule sets being implemented are typically considered company secrets by MSS providers—partly because of the risk that an attacker may abuse this information to attack customers of the MSS, and partly because some rules may implement possibly privacy invasive IT monitoring policies, for example monitoring use of peer-to-peer traffic, if the IT policy disallows this. The network owner may not want to reveal such monitoring practices, since such information would be considered sensitive from a business or reputation perspective.

An important principle should be that the monitoring invades privacy and confidentiality as little as possible for normal, unsuspicious traffic. However, just like in the airport example, a more thorough investigation will be required if suspicious Internet traffic is detected, to verify whether the data traffic is hostile or not.

The decisions and actions security analysts perform should be logged, regardless of whether the analyst decides to investigate an event in detail or not, since this will provide transparency on what is being investigated and why, both from a security and privacy perspective. Such transparency can be expected to be instrumental in improving the MSS operation both from a work efficiency, attack detection and privacy impact perspective, since it would allow identifying and putting effort into mitigating bottlenecks, blind spots or overly privacy invasive sides of the operation.

[1] There will also be synergies between privacy enhancing technologies and security, as will be discussed later. Aiming for such synergies is recommended by the 4. Privacy by Default principle, which states that one should aim for a win-win situation between privacy and security (Cavoukian et al. 2010).

Information Leakages from Intrusion Detection and Prevention Systems

The first question one perhaps should ask, is whether there really is any significant leakage of private or confidential information in IDS alarms from outsourced MSS? The market leaders claim that false alarms is not a problem for a properly managed IDS/IPS in their technical documentation. However, a recent comparative analysis of commercial IDS and IPS indicates that more than 92.85 % of all IDS alarms on a campus network from a test bed of seven different commercial IPS/ IDS products, tested over a period of 2 years, are false alarms (Ho et al. 2012). Other studies have also indicated that network monitoring technologies create a significant amount of false alarms (Alharby and Imai 2005). Some of the reason for this, is applications that do not follow the protocol specifications (Ho et al. 2012).

Furthermore, around 91 % of the false alarms were not related to security issues, but management policies, for example that IDS rules were set up to identify peer-to-peer (P2P) traffic, that was not allowed according to the IT policy (Ho et al. 2012). One such example is the Snort rule *sid:1427* "MULTIMEDIA Windows Media download", which is a broad IDS rule that matches download of any windows media files via the web. This IDS rule is shown below:

```
alert tcp $EXTERNAL_NET 80 -> $HOME_NET any (
msg:"MULTIMEDIA Windows Media download";
flow:from_server,established;
content:"Content-Type|3A|"; nocase;
pcre:"/^Content-Type\x3a\s*(?=[av])
      (video\/x\-ms\-(w[vm]x|asf)|
      a(udio\/x\-ms\-w(m[av]|ax)|
      pplication\/x\-ms\-wm[zd]))/smi";
classtype:policy-violation;
sid:1437;
rev:6;)
```

This IDS rule matches any HTTP response messages originating from the external network and with destination

to any port on the home network. The rule triggers on a case-insensitive regular expression on content matching the Windows multimedia MIME types for *wvx, wmx, wma, wmv, wax, wmz* and *wmd* files. The rule has class type policy-violation which broadly means a violation of a corporate IT policy. There are in total 720 such policy-violation rules in the Snort ruleset at time of writing. These policy rules are disabled by default in Snort, which is commendable, however the mere presence of such rules cause concern and uncertainty, since they may cause a significant leakage of private or confidential information if enabled. It is furthermore typically not possible to reveal whether such IDS rules are enabled or not, since this is considered business sensitive confidential information by the MSS providers. Another problem is that Internet applications do not follow the standards, and therefore may trigger false alarms, for example on web-based IDS rules that check for standards conformity (Ho et al. 2012).

A similar area is IDS rules for identifying web bugs. Web bugs are objects that are embedded in a web page or email that usually is invisible to the user, but allows for checking whether a user has viewed the web page or email. They can for example be implemented using one pixel transparent GIF images or using Javascript.[2] These web bugs may in themselves be a risk for privacy and data confidentiality, since they track who is reading an email or web page when, and from where. However, in this case, the good intention of security monitoring may be its own worst enemy, because detecting privacy leaking web bugs using intrusion detection systems may cause a significant privacy leakage in itself. The reason for this is that IDS rules trigging on web bug activity also may trigger on user sessions referencing these web bugs. However, in this case it should only be interesting from a security perspective to detect the presence of web bugs. It should not in general be necessary to view possibly privacy-leaking payload, addresses or subsequent advertisements that these plugins may cause, perhaps apart from a limited analysis of the effect these web bugs have from a security or privacy perspective.

It is in other words not difficult to show that intrusion detection systems can be configured to leak information. The evidence should rather be on computer security companies to prove or certify that they do not violate privacy to a greater extent than necessary when using current state-of-the-art technology. The current large void on what indeed is being monitored by computer security companies is the big problem. Snort has for example 4,844 IDS rules enabled by default (including proprietary binary rules) of a palette of over 21,000 IDS rules. Similar numbers can be expected for other commercial actors that use IDS rules based on attack signatures. This is especially problematic with respect to the large amount of possibly privacy-leaking policy rules, that may or may not be enabled depending on corporate IT policies—we neither know whether they are being used or not, nor what they trigger on. One way to elucidate this problem, would be to develop a metric that is able to benchmark how privacy invasive a MSS operator is. This would allow MSS operators to compete on merit, both from a privacy and security perspective.

Furthermore, transparency and nonrepudiation is important for ensuring auditability of activities on private or confidential information, which means that secure logging schemes are required to be able to prove who have accessed the given information when. Anonymisation, pseudonymisation or encryption are general techniques that can be used to reduce the privacy impact in cases like this, given that suitable metrics exist for identifying *where, what* and *how much* sensitive information that leaks. Sensitive parts of the data should then be anonymised, possibly using a reversible anonymisation scheme based on a combination of anonymisation and encryption, so that only authorised stakeholders can access this information. At the same time, access to such information should be logged to ensure transparency and accountability.

Another risk is that the monitoring organisation may not act morally right and abuse acquired knowledge from private or confidential information. Corrupt insiders in the monitoring organisation may for example sell private or confidential information, extort the information owner or use the information for their own advantage (Radianti and Ulltveit-Moe 2008). A more recent concern is the risk of radicalisation by insiders that have access to private or confidential information, especially for critical infrastructures. One important principle here, is that the monitoring organisations must be accountable and auditable for the operations they do on private or confidential information. This requires that techniques for ensuring *transparency* and *non-repudiation* are built into the monitoring technologies, so that the monitoring organisation cannot deny having processed given private or confidential information.

A recent trend that is expected to be the next major advance in attack detection, is to merge Big Data analytics with IDS/IPS, so that all communication to or from a company can be stored and investigated over a time span of months. An early example of this is the time machine (Maier et al. 2008), which works in a similar way as a "Personal Video Recorder" for network traffic, being able to store the initial part of all network sessions. Now more powerful cluster-based technologies like Apache Hadoop have been combined with IDS technologies to log all network traffic in real time, and at the same time perform near real-time attack analysis on this traffic. An example of this is PacketPIG,

[2] Web bug definition: http://en.wikipedia.org/wiki/Web_bug

which is capable of storing all data from a 100Mbit/s link in real-time for months on a 3Tb disk (Baker et al. 2012).

An advantage with such technologies, is that they allow detecting some formerly unknown attacks (so-called zero-day attacks) after the attack has happened, by performing a retrospective IDS analysis on stored data. This approach works under the assumption that the IDS is rechecking the stored data using new attack signatures identified after the data was logged. This improves the capabilities for performing data forensics significantly. However, these techniques also cause a concern both from a privacy and transparency perspective, since the operation on such big data is concealed in legal and commercial secrecy (Richards and King 2013). Another problem is that these techniques normally detect and not deter attacks, since they are based on data mining of past traffic. There is furthermore a lack of mechanisms for protecting the privacy and confidentiality when accessing these big data, as well as lack of logging mechanisms for ensuring transparency and non-reputability. Much of the reason for this, is that big data based security analysis still is in its infancy.

Privacy and security interests compete on funding

Privacy-enhancing technologies should be used to allow security monitoring being performed as precisely as possible, in order to minimise the privacy and confidentiality impact of MSS operations. A challenge is that security monitoring needs to be implemented within a commercial organisation that mainly aims to maximise the profit for its owners. This means that a customer of a MSS provider will have a limited budget available for security investments. It has for example been suggested that only a fraction of the expected loss due to security breaches (max 37 %) should be spent on security investments for a risk neutral firm (Gordon and Loeb 2002). This also means that privacy and security interests need to compete on the funding to implement the best possible security and privacy handling.

There are in other words practical limits for how much money and effort that a monitoring company should put into both security and privacy to improve the service. This means that solutions for enhancing the privacy should be readily available, affordable and easy to configure, preferably over existing services, to reduce the implementation costs for adding privacy and confidentiality protection. There is otherwise a risk that security interests may trump the privacy interests given a limited budget. This is at the moment a major hurdle, since technologies for privacy-enhanced security monitoring are not yet readily available. Later a roadmap on how this deficiency can be mitigated will be presented.

The monitoring organisation may also see synergies from better privacy handling from an economic perspective, for example if improved handling of private or confidential information has side effects like reduced operating costs from handling fewer false alarms, or better protection of corporate secrets or personally identifiable information. This is in-line with the 4. foundational Privacy by Design principle (Cavoukian et al. 2010), since integrating privacy enhancing technologies should aim at creating a win-win situation by supporting both the privacy and security objectives. In addition, improved privacy handling reduces the risk of liabilities from privacy leakages, and it will improve the trustworthiness for customers where privacy and confidentiality is paramount. One example of such customers is health institutions who, due to very strict privacy requirements, will not allow sensitive data to leave the corporate network.

Measuring privacy

A Utilitarian way to describe the optimal utility level has been proposed based on information theory (Sankar et al. 2010):

> "For a data source with private and public data and desired utility level, maximum privacy for the private data is achieved by minimising the information disclosure rate sufficiently to satisfy the desired utility for the public data."

This implies that private or confidential information is disseminated strictly on *a need to know* basis. An advantage with this approach, is that it may be possible to quantitatively analyse the optimal solution and compare how close a real solution is to the optimal one, given that some objective criteria or *metrics* for the information disclosure rate are identified. Privacy metrics like differential privacy have been proposed as a method to quantify the maximum privacy for a given level of utility for cases where sensitive data in databases need to be sanitised, for example by adding noise to blur the precision of given data, while still maintaining important statistical qualities, like the mean and standard deviation over a sufficiently large sample (Dwork 2006).

A disadvantage with this model is that it does not consider the semantics and therefore not the *value* of revealed private data. Some data are typically considered more sensitive and therefore also more valuable than other. Econometric or actuarial models have been suggested for modelling the cost of revealing data (Gritzalis et al. 2007; Yannacopoulos et al. 2008). The practical challenge with these economic models, is that it may be difficult to get representative cost distributions, since they are based on people's subjective value of private data.

Estimating the value of privacy

I did some preliminary experiments as part of my research where security analysts attempted to classify the privacy leakage of IDS alarms. They found it very difficult to do this. In many cases they found it hard to understand, or even purely hypothetical, that the sampled IDS alarms even would contain any significant information that was sensitive from a privacy or confidentiality perspective. The information they sampled, was after all open (i.e. not encrypted/protected). This illustrates the well-known fact that the perceived value of private information is highly subjective (Gritzalis et al. 2007), so it is not given that the valuation by security analysts, which are the only people that have security clearance to access the IDS alarms, would give a representative picture of the privacy leakage. In practice, the only stakeholders that can give the correct valuation of the privacy impact from IDS alarms, are the users themselves. And it is in most cases not trivial and also not desirable to connect the users to the underlying data from a privacy perspective.

One possible way to get around this problem, to get realistic measurements of the privacy impact may be the following: Assume that the data controller compiles a top ten list of the most privacy concerning information leakages, for example from a given web service. The data controller then needs to ask a representative random sample of users in an anonymous poll, presented during use of the service, what they think their privacy is worth in monetary value, given that a security company may see how they used a given set of web pages. The results from this poll could then be used to estimate a privacy impact factor as a random variable for each given information leakage.

It may however in practice not be feasible to do this, because it would be difficult to get permission to do such an experiment in an outsourced scenario where you would have to consider the business concerns of both the MSS provider and the service provider being monitored. It is hard enough to get consent from the MSS provider to do research on IDS data, and may be even harder to get consent from customers of MSS services, due to concerns that such a detailed poll would affect the reputation of the service being monitored. This means that it will be challenging at best, maybe not even possible, to get a representative cost distribution for the privacy impact of the information leakages, not to mention getting a representative cost distribution for an entire IDS rule set consisting of several thousand rules. Furthermore, privacy valuation is very sensitive to how the question is framed (Acquisti et al. 2010).

Another challenge, is that the value of private data changes over time, and may either increase or decrease (Berthold and Böhme 2010). Privacy valuation has for example been investigated based on option pricing theory,

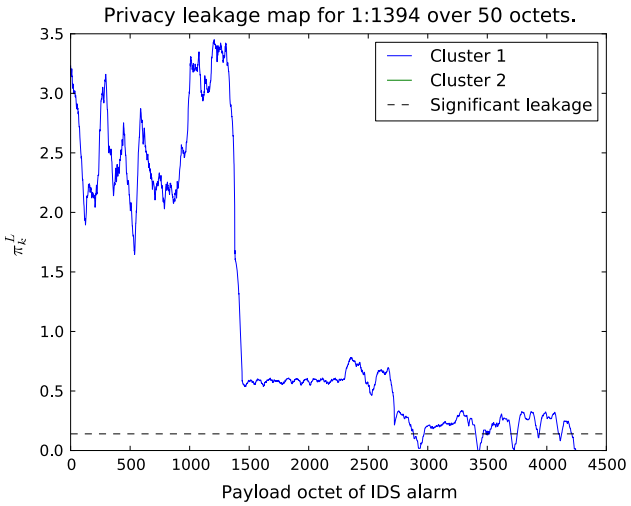

Fig. 1 Privacy leakage map for Snort IDS rule 1:1394 detecting NOP sled based buffer overflow attacks

where the self-information of a private data item is simulated over time using a stochastic random walk (Berthold and Böhme 2010). It is however hard to predict whether the value of private information will increase or decrease in value over time, except in trivial cases. One such example is linkability between targeted advertisements (e.g. from doubleclick.net). Such targeted advertisements may be problematic from a privacy perspective, since they may reveal personal preferences, however these advertisements also typically time out after a relatively short period, meaning that the information after this becomes worthless.

How to measure privacy leakage

The discussion in the previous section indicates that is is better to focus on measuring information leakage in IDS alarms based on objective criteria which correlate with the disclosure rate of sensitive information, for example based on Shannon entropy (Shannon 1948), rather than doing detailed privacy valuation analysis.

A privacy leakage metric for IDS, founded on the theory of quantitative information flow analysis (Smith 2009, 2011) and Shannon entropy is proposed in (Ulltveit-Moe and Oleshchuk 2013). This metric is from a high-level perspective based on measuring the entropy standard deviation around each attack vector cluster an IDS rule appears to match. Qualitative analysis of attack vector behaviour and quantitative simulations have verified that this could be a good model of information leakage from IDS rules.[3]

[3] The detailed theory behind this metric is considered beyond the scope of this paper, but interested readers can read the full paper here (Ulltveit-Moe and Oleshchuk 2013).

Figure 1 illustrates how this privacy leakage metric can be used for analysing where information leaks in the payload of an IDS alarm are, as well as analysing what this information means. The privacy leakage map is based on IDS alarms from my own test network and are only valid for the given measurement data. The figure does however illustrate how the concept for measuring privacy leakage works. The IDS alarm with SID 1:1394 "Shellcode x86 NOP" aims at detecting the NOP sled, which malware uses to overwrite the stack of the attacked system, in order to exploit a vulnerability. This IDS rule is very simple, and triggers on a sequence of at least 31 'A' characters. It is well known that this IDS rule is overly broad and has a large number of false alarms.[4] This rule shows significant information leakage up to around 1,400 octets, where it drops step-wise off. Investigations of the data shows that data cluster 1 triggers on random web traffic, for example HTTP response and Set-cookie: messages from various web sites. This is obviously false alarms, and is problematic from a privacy perspective, since it reveals user behaviour. Cluster 2 matches packets with only 'A' characters. These have zero variance and are least concerning from a privacy perspective, since they precisely match the attack definition in the IDS rule. This illustrates a type of broad IDS rules that have poor attack detection capabilities and therefore may create a lot of false alarms. This and similar IDS rules would benefit from improving the attack detection pattern to be more specific, which would both reduce the number of false alarms as well as reduce the information leakage from the IDS rule, thereby improving the rule from a privacy perspective. This IDS rule is disabled in newer versions of the Snort ruleset, since it has low utility from a security perspective. It can also be observed that another way to reduce the information leakage is to encrypt or anonymise the information, which would reduce entropy variance and therefore also the measured privacy leakage to an insignificant level (Ulltveit-Moe and Oleshchuk 2013).

This privacy leakage metric furthermore allows a data controller to set an impact value on given data, based on the perceived sensitivity of the data. This is not an exact valuation of the private data, however it can be useful to weigh up data that clearly is more sensitive from a privacy or confidentiality perspective, for example confidential or graded information. It can also be used to reduce the impact of data that by investigation clearly has no or little value from a privacy or confidentiality perspective. This is a simplistic approach similar to what is common in risk analysis. The metric furthermore uses Expectation Maximisation-based clustering, based on the solutions in (Cord et al. 2006; Figueiredo and Jain 2002), for identifying attack vector clusters that the IDS alarms trigger on (Ulltveit-Moe and Oleshchuk 2013). This is a metric that can be used for measuring privacy leakage in IDS alarms. The metric is also useful in other scenarios, for example for verifying correct enforcement of a privacy policy, since an effective privacy policy will cause a reduction in measured entropy standard deviation for the data elements the anonymisation policy operates on.

[4] See: http://www.snort.org/search/sid/1394

Attack prevention or surveillance, which is better?

Intrusion Prevention Systems (IPS) is a network monitoring technology that extends IDS with the possibility to automatically enforce a computer security policy. A question is then: when is it acceptable from an ethical/moral perspective to automatically enforce a computer security policy, and are there any cases where it can be considered better to automatically enforce the policy than to use traditional monitoring techniques like IDS? A related question is whether blocking of undesirable content is more acceptable than surveillance covering use of undesirable content?

In general, IPS, firewalls and IDS may all leave electronic evidence in the form of system logs or alerts sent to a central security operations centre. It is possible to define rules that enforce a security policy without leaving electronic traces, however this is not common to do. The reason is that system logs are useful to detect and improve rules that perform poorly or incorrectly. It can also be useful to verify correct system operation, as well as for correlating different alerts in order to infer attacks with greater confidence.

Logging of what is being monitored may also be important for accountability, to audit what is being monitored either by the network owner or by third party quality certification organisations. It should however be noted that such logs also may contain private or confidential information. They should therefore be cryptographically protected both against unauthorised modifications by the MSS provider as well as against external attacks, and should use privacy enhancing technologies, for example anonymisation or pseudonymisation, to avoid showing private or confidential information in cleartext to unauthorised personnel. Since IPS rules typically perform automated actions, then there should normally not be a need to view detailed information from such events in cleartext. IPS alerts should therefore be suitable candidates for anonymisation or pseudonymisation.

A problem with automatic enforcement using IPS, is however that monitoring rules typically are neither perfect, meaning that false alarms may occur, nor complete, meaning that the rule is not able to catch all attacks (Flegel 2007). This means that using an IPS causes a risk that some

legitimate traffic also will be denied. On the other hand, one should not be complacent because of having an IPS implemented, since the rule definitions typically are not complete, and may not detect all attack scenarios.

A common way for IPSs to enforce preventive actions, is to block traffic from the attacker either permanently or for a given time interval. This can be problematic both from an ethical and business perspective since it may cause benign traffic to be blocked out. There is also a risk of targeted Denial of Service attacks against the IPS or firewall if the adversary uses forged attack traffic to disrepudiate a given user or to block the entire service. This shows that automatic filtering blocking attack traffic that matches given rules can be problematic from both an ethical, business and security perspective, although it clearly is more cost effective than manual 24x7 monitoring of IDS alerts. It is also more efficient since it actually may prevent an ongoing attack, given that the IPS is sufficiently fast and precise to detect and deter the attack vector without harming innocent third parties.

A somewhat related area, is permanent blacklisting of traffic from certain hosts assumed under control by adversaries, or even censorship of web sites providing content that in a given legislation is deemed illegal. Is automatic enforcement of security policies, for example via rules that deny access to certain on-line resources in this case more acceptable than security monitoring? Is it for example worse to block inappropriate web sites or web sites that may be risky from a security perspective, than if humans investigate such events? This is a discussion on censorship versus surveillance—which one is better or worse. Content filtering is cheaper and may be a better choice from a purely economical perspective, however one may risk liabilities from legitimate users and customers whose service has been interrupted. The other extreme, is whitelisting where only traffic between approved entities is allowed. Such approaches may be useful in certain scenarios, for example for controlling access to critical infrastructures.

Content filtering can be considered better from a privacy perspective provided that IPS alarms are properly anonymised. However it is not necessarily better from an anti-censorship/free speech perspective. Knowing that systems in general log what is being filtered means that it can be discussed whether content filtering is a good argument from a privacy perspective, although it certainly is possible to create IPS rules that either anonymise or encrypt sensitive information or do not log any information at all. Also, a censored environment may give a deceptive perception of reality, something that is morally questionable.

Content filtering using IPS or firewall technologies is in other words useful and can be morally acceptable if used against attack scenarios, provided that the MSS provider aims at minimising the harm from both a privacy and freedom of speech perspective, as well as avoiding harm for innocent third parties. However a potential risk is that the IPS may be vulnerable to denial of service attacks.

A roadmap towards improved privacy for managed security services

The roadmap describes what is needed to reduce the privacy problems of outsourced monitoring using intrusion detection systems. There will need to be interaction between different stakeholders, public authorities and policy makers in order to improve privacy for managed security services.

Figure 2 shows a roadmap for how privacy-enhanced intrusion detection systems can be made reality. The approach for achieving this is based on a technology development process. It starts with policymakers funding research projects that subsequently can implement the necessary technologies and methodology in several iterations. Early adopters, for example critical infrastructures and hospitals, can then start using the technology, and their experiences from using privacy-enhanced IDS can be used as input to a standardised measurement methodology with supporting technologies for measuring and reducing privacy leakage. When technology and appropriate standards are in place, then policymakers can consider amending existing legislation to favour or require privacy-enhanced operation. This in turn would lead to commercialisation and general deployment of privacy-enhanced IDS services. Details on how this can be achieved is discussed in the subsequent subsections.

Research programmes for privacy-enhanced surveillance techniques

Research programmes supporting the development of privacy-enhanced IDS are needed to develop the necessary technical and methodological foundations. Policymakers in both EU and US have announced research programmes that fund research on such privacy enhancing technologies. The National Science Foundation for example funds research on policy driven frameworks for online privacy protection and privacy-aware information release control which may be relevant for privacy enhanced intrusion detection systems (National Science Foundation 2014). The EU seventh research programme had strong focus on privacy in the Security programme, and the Horizon 2020 programme extends this by supporting privacy and security as cross-disciplinary concerns (European Communities 2014). There is in other words a significant amount of research funding available for research on privacy-enhancing

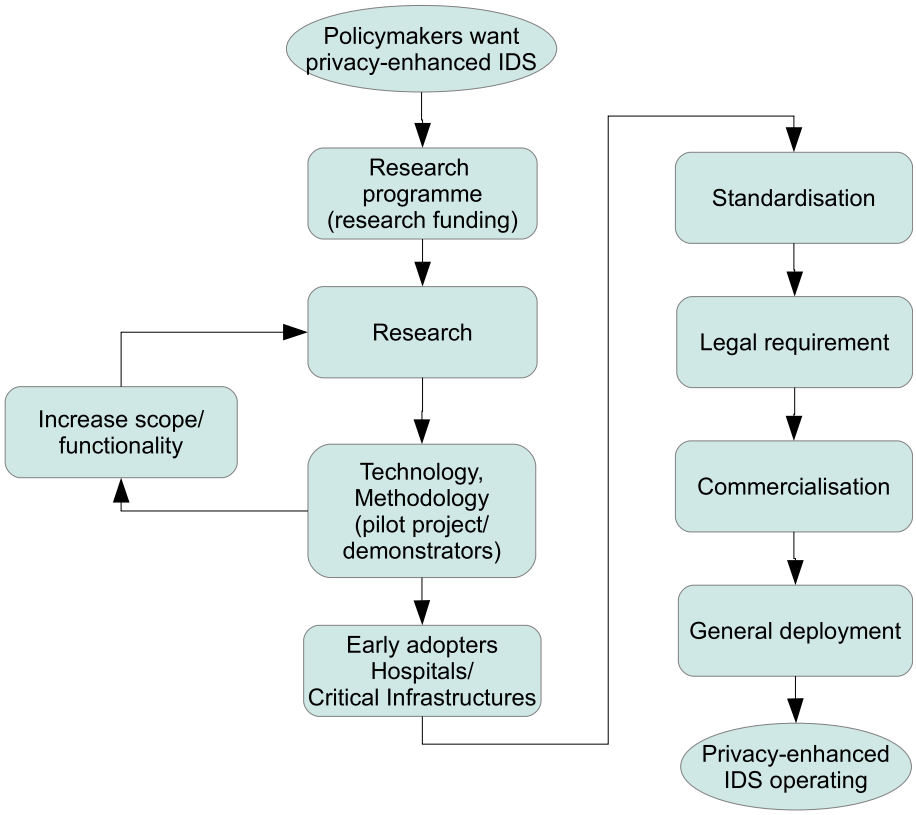

Fig. 2 High-level roadmap for implementing privacy-enhanced intrusion detection services

technologies, showing that policymakers have commitment for making such services a reality.

Research on privacy-enhanced security operations

Privacy-enhanced IDS was a popular research area some years ago, however it seems to have lost some momentum. Some of the reason for this may be that it is difficult to do research on IDS due to lack of realistic data sets (Tavallaee et al. 2009). It is furthermore very difficult to do research on real security operations, since these typically require high security vetting. It may also be a challenge to publish research findings, since these first will need to be checked and declassified before they can be published.

Currently a limitation is that technologies for supporting privacy-enhanced operation and suitable privacy metrics are not readily available. This means that technology investment costs initially will be high for such projects, and the incentive for using such technologies low, since existing legislation in most cases does not mandate using them. This can be mitigated by developing the necessary privacy-enhancing technologies as part of research projects. Such research projects could develop the necessary privacy leakage metrics and measurement methodology based on existing best practices like privacy impact assessments and the Privacy by Design requirements. It is important that this research provides solutions that are efficient and usable and can be integrated with or retrofitted on existing monitoring technologies, in order to lower the threshold and costs for adopting these new technologies.

Privacy leakage metrics are needed for quantifying *how much* private or confidential information that is leaking, *where* these leakages occur and for identifying *what* this information means. This is required to support a continuous improvement process according to the well-known Plan Do Check Act method of improvement (Moen et al. 1999), in order to gradually improve privacy protection of security monitoring techniques over time. Research on suitable metrics for detecting the primary privacy leakages should, due to the lack of realistic test data, initially be based on theoretical models verified via simulations as a first step towards making privacy leakages measurable (Ulltveit-Moe and Oleshchuk 2013). It may at a later stage be possible to verify how well these models work in a realistic environment when privacy-enhanced IDS has been commercialised and more widely deployed. This essentially means defining a gold standard for privacy leakage measurements, that at a later point could be modified or adjusted if research shows that it has problematic biases. We have suggested one such gold standard based on the

standard deviation of Shannon entropy (Ulltveit-Moe and Oleshchuk 2013).

The research on suitable metrics for detecting privacy leakages would need to be paralleled by research on privacy-aware information release control, so that only authorised stakeholders can access sensitive information during security operations while at the same time ensuring transparency, accountability and nonrepudability for such transactions.

Management of privacy policies should be done using an overarching methodology. Such policies will typically be defined after a privacy and security impact assessment, which identifies the data elements and services that need to be protected. The privacy impact assessment may also identify under which conditions data are sensitive, for example if an identifiable subset of IDS alarms should be anonymised. The privacy policy should also support giving different stakeholders access to different parts of the private or confidential information according to their needs. This means that the privacy enforcement solution should support multi-level security. It should also allow defining both default PERMIT privacy policies, where any information that is explicitly being authorised may be anonymised, and default DENY policies, which by default anonymise all information, and where selected information deemed safe from a privacy or confidentiality perspective subsequently can be declassified. The latter allows for supporting Privacy by Default according to the Privacy by Design principles (Cavoukian et al. 2010).

The technologies and methodology for privacy enhanced intrusion detection produced by research projects should subsequently be demonstrated and tested in realistic conditions, in order to ensure that both the technologies and methodology work as expected.

Iterations to increase scope and functionality

The research should start by implementing the core functionality required to achieve privacy-enhanced IDS. When this functionality works, then other functions can be added, for example:

- Adding privacy enhancing technologies to digital forensic interfaces;
- Secure logging to support transparency and nonrepudiation of actions and transactions from privacy-enhanced IDS;
- Improving performance of the privacy enhancing technologies, to increase the alarm handling capacity of anonymisers and deanonymisers and reduce latency;
- Privacy-enhanced alarm correlation systems, effectively implementing higher-order intrusion detection systems;
- Protecting information about threats exchanged between partner organisations. This allows sensitive threat information to be anonymised, however a trusted service may be allowed to deanonymise and operate on certain parts of this information;
- Private or confidential information can furthermore be protected using homomorphic encryption or similar techniques, to allow certain operations on encrypted data in a similar way as CryptDB does (Popa et al. 2011);
- Adding privacy protection to big data analytics based IDS and data forensics solutions;
- Furthermore, privacy leakage measurements should consider the anonymity set of the underlying data using metrics like k-anonymity or l-diversity (Sweeney 2002; Ciriani et al. 2007; Machanavajjhala et al. 2007), in order to reduce the risk that data mining based on publicly known information could be used to deanonymise sensitive data.

Achieving a reasonable coverage of all this functionality would be a long term goal. The short term goal would be to start by implementing the core functionality (privacy-enhanced alarm handling) first, and then extend the technology to cover other use cases later. This means that the technology and methodology should be designed to be general and extensible.

Early adopters of privacy-enhanced IDS

As the technology matures, then it will be possible for first movers, for example within the health sector or critical infrastructures, to try out the privacy-enhanced technologies and techniques as illustrated in Fig. 3. First movers should be able to try out the technologies 1–3 years after the research projects have concluded. This could also spark commercial interest for the privacy-enhanced technology by privacy-conscious MSS operators that want to market improved handling of private and confidential information as an added value. Having the privacy-enhanced techniques freely available as Open Source software will lower the threshold for adopting such services, and may also be used to improve the tools and techniques as a community effort.

It is also important to increase public awareness about the privacy-enhancing technologies, once they are ready for market, in order to capitalise on the public opinion's scepticism against surveillance techniques. It would on the other hand also be necessary to work on reducing the scepticism and barriers for adopting privacy-enhancing technologies by commercial managed security service providers. The latter can probably best be mitigated by providing a solution that is easy to use and which integrates well with existing systems. Such a solution should also

Fig. 3 Critical infrastructures like the health and transport sector could be early adopters that may benefit from using privacy-enhanced intrusion detection systems

demonstrate the possible synergies that can be achieved between privacy and security objectives, in order to convince MSS providers that there is a real benefit—also from a security perspective from using such technologies. There are already legal incentives for using such technologies in some areas, for example for health institutions or critical infrastructures which by definition handle confidential information. This means that there already probably is a market for privacy-enhanced intrusion detection systems, although it currently is not very large.

Standardisation

The technologies and methodology should be standardised once good technical solutions have been identified by the research projects and early adopters. Standardisation should include templates for incentive-compatible service level agreements that consider both the security and privacy/confidentiality side of the operation. To increase the adoption rate, the standards should preferably be open and available at no cost, to reduce the risk that privacy-enhanced technologies are not being adopted by vendors for cost or IPR reasons. Research projects can provide the inception for technologies that later can be standardised by standardisation organisations in order to provide both technologies and privacy leakage measurement methodologies that are open, repeatable and consistent over time. Such standardisation activities may take long time, depending on how conflicting the interests between different commercial vendors are.

Legal requirements

When the technology and measurement methodology is in place, then policymakers should support the technological advancements in privacy enhancing technologies for monitoring technologies with legal regulations or directives that mandate use of such technologies. This means that transparency and nonrepudability of actions on private or confidential information can be supported by outsourced Security Operations Centres. It will necessary take time (several years) until such regulations can be enforced, since the security industry will need time to adapt to the new legal requirements.

General deployment

Once the technology has matured and has been standardised and incorporated into commercial products, then it will be possible to deploy privacy-enhanced intrusion detection systems on a larger scale, as mandated by updated laws and regulations. This is when the general public can be expected to benefit from such privacy-enhanced intrusion detection services.

Conclusion

This paper does an analysis of ethical and privacy issues related to outsourced managed security services based on intrusion detection systems. The analysis shows that there may be a significant risk of such systems leaking private or confidential information, especially in scenarios where managed security services have been outsourced. The lack of quantitative metrics for identifying privacy leakage in such systems, together with the veil of secrecy due to operating on information that is graded, means that it currently is not possible to quantify this privacy leakage risk in a repeatable, comparable and coherent way. This means that it currently is not possible to compare managed security services on how good they are from a privacy perspective, it is only possible to quantify the security side

of the operation in the form of number of detected attacks, and possibly also the false alarm rate.

An effect of this, is that privacy currently is considered a non-issue in commercial intrusion detection systems, unless such technologies are used in sectors where extra precaution is needed, for example for health institutions or critical infrastructures. In these cases privacy handling is attempted introduced as an afterthought, for example by using techniques like thin clients for monitoring IDS data inside the hospital or critical infrastructure in an attempt to avoid that sensitive data leaves the hospital perimeter. This strategy has obvious limitations since it lacks transparency and accountability on who have accessed which private or confidential information when. The current technology also lacks privacy enforcement mechanisms and metrics for identifying and limiting leakage of private or confidential information over time.

The paper has proposed a roadmap on how these issues can be mitigated by developing suitable privacy enforcement mechanisms in combination with a gold standard for privacy leakage measurements that is able to quantify how much information that is leaking, where these leakages are and what these information leakages mean. This gives the data controller a much more fine-grained mechanism for measuring and controlling privacy leakages, so that they can be reduced to an acceptable level over time. The privacy leakage metric would allow comparing managed security service providers on equal terms from a privacy perspective without revealing any sensitive operational details. The roadmap shows that it should be feasible to implement the necessary technologies and methodology required to enforce privacy control for managed security services. If the privacy enforcement mechanisms and privacy leakage metrics subsequently are standardised, then this opens up for stronger privacy protection also from a legal perspective in the future, since managed security services then would be auditable to a much larger extent than they are today. This would allow managed security service providers to compete on merit both from a privacy and security perspective.

Acknowledgments Thanks to all anonymous reviewers, for challenging questions and good ideas on how to improve the quality of the paper. This work has been partially supported by the project "PRECYSE - Protection, prevention and reaction to cyber-attacks to critical infrastructures", funded by the European Commission under the FP7 frame programme with contract number FP7-SEC-2012-1-285181 (www.precyse.eu), and partially by Telenor Research and Innovation under the contract DR-2009-1.

References

Acquisti, A., John, L., & Loewenstein, G. (2010). What is privacy worth? http://www.futureofprivacy.org/wp-content/uploads/2010/07/privacy-worth-acquisti-FPF.

Alharby, A., & Imai, H. (2005). IDS false alarm reduction using continuous and discontinuous patterns. *Lecture Notes in Computer Science*, *3531*, 192–205.

Baker, M., Turnbull, D., & Kaszuba, G. (2012). Finding needles in haystacks (the size of countries). http://media.blackhat.com/bh-eu-12/Baker/bh-eu-12-Baker-Needles_Haystacks-WP.

Berthold, S., & Böhme, R. (2010). Valuating privacy with option pricing theory. In T. Moore, D. Pym, & C. Ioannidis (Eds.), *Economics of Information Security and Privacy* (pp. 187–209). US: Springer.

Bicknell, P., & Jean, H. (2011). National information assurance partnership common criteria evaluation and validation scheme, validation report hp tippingpoint intrusion prevention systems. http://www.commoncriteriaportal.org/files/epfiles/st_vid10345-vr.

Cavoukian, A. (2009). Whole body imaging in airport scanners: Activate privacy filters to achieve security and privacy. http://www.ipc.on.ca/images/Resources/wholebodyimaging.

Cavoukian, A., Taylor, S., & Abrams, M. E. (2010). Privacy by design—Essential for organizational accountability and strong business practices. *Identity in the Information Society*, *3*(2), 405–413.

Ciriani, V., di Vimercati, S. C., Foresti, S., & Samarati, P. (2007). k-Anonymity. In: Secure data management in decentralized systems (pp. 323–353). Springer.

Cisco (2013). Writing custom signatures for the cisco intrusion prevention system. http://www.cisco.com/web/about/security/intelligence/ips_custom_sigs_pdf.

Cord, A., Ambroise, C., & Cocquerez, J. P. (2006). Feature selection in robust clustering based on laplace mixture. *Pattern Recognition Letters*, *27*(6), 627–635. doi:10.1016/j.patrec.2005.09.028.

Ding, W., Yurcik, W., & Yin, X. (2005). Outsourcing internet security: Economic analysis of incentives for managed security service providers. In: Internet and network economics, LNCS, vol 3828 (pp. 947–958). Springer.

Dwork, C. (2006). Differential privacy. Automata, languages and programming (pp. 1–12).

European Commission. (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:NOT.

European Communities. (2014). Digital security: Cybersecurity, privacy and trust. URL http://ec.europa.eu/research/participants/portal/desktop/en/opportunities/h2020/topics/99-ds-01-2014.html.

Figueiredo, M. A. T., & Jain, A. K. (2002). Unsupervised learning of finite mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *24*(3), 381–396. doi:10.1109/34.990138.

Flegel, U. (2007). *Privacy-respecting intrusion detection* (1st ed.). Berlin: Springer.

Gartner. (2010). Magic quadrant for network intrusion prevention systems. URL http://mcafee.zinfi.com/enduser/ngns/dyntek1/file/McAfee_vol4-art5.

Gordon, L. A., & Loeb, M. P. (2002). The economics of information security investment. *ACM Transactions on Information and System Security*, *5*(4), 438–457. doi:10.1145/581271.581274.

Gritzalis, S., Yannacopoulos, A., Lambrinoudakis, C., Hatzopoulos, P., & Katsikas, S. (2007). A probabilistic model for optimal insurance contracts against security risks and privacy violation in IT outsourcing environments. *International Journal of Information Security*, *6*(4), 197–211. doi:10.1007/s10207-006-0010-x.

Ho, C. Y., Lai, Y. C., Chen, I. W., Wang, F. Y., & Tai, W. H. (2012). Statistical analysis of false positives and false negatives from real traffic with intrusion detection/prevention systems. *IEEE*

Communications Magazine, *50*(3), 146–154. doi:10.1109/MCOM.2012.6163595.

Kairab, S. (2005). *A practical guide to security assessments*. Boca Raton, Florida: Auerbach Publications.

Richards, N. M., King, J. H. (2013). Three paradoxes of big data. Stanford Law Review Online 66:41, URL http://www.stanfordlawreview.org/online/privacy-and-big-data/three-paradoxes-big-data.

Machanavajjhala, A., Kifer, D., Gehrke, J., Venkitasubramaniam, M. (2007). l-diversity: Privacy beyond k-anonymity. Cornell University p 52, URL http://www.truststc.org/pubs/465.html.

Maier, G., Sommer, R., Dreger, H., Feldmann, A., Paxson, V., & Schneider, F. S. (2008). Enriching network security analysis with time travel. *SIGCOMM Computer Communication Review*, *38*(4), 183–194. doi:10.1145/1402946.1402980.

McAfee. (2007). Mcafee intrushield IPS, user-defined signature creation version 4.1. https://kc.mcafee.com/resources/sites/MCAFEE/content/live/PRODUCT_DOCUMENTATION/20000/PD20345/en_US/INTR_User-Defined_Signatures_4.1.

Moen, R. D., Nolan, T. W., & Provost, L. P. (1999). *Quality improvement through planned experimentation*. New York: McGraw-Hill.

National Science Foundation. (2014). US NSF-CISE-funding. URL http://www.nsf.gov/cise/funding/cyber_awards.jsp.

Plashchynski, D. (2014). viewssld—SSL traffic descryption daemon. URL http://sourceforge.net/projects/viewssld/.

Popa, R. A., Redfield, C. M. S., Zeldovich, N., & Balakrishnan, H. (2011). CryptDB: Protecting confidentiality with encrypted query processing. In: Proceedings of the twenty-third ACM symposium on operating systems principles, ACM, New York, NY, USA, SOSP '11, (pp. 85–100), doi:10.1145/2043556.2043566.

Radianti, J., & Ulltveit-Moe, N. (2008). Classification of malicious tools in underground markets for vulnerabilities. *NISK*, *2008*, 19–31.

Roesch, M. & Green, S. C. (2009). Snort. URL http://www.snort.org/assets/82/snort_manual.

Samuelson, P. (2000). Privacy as intellectual property? Stanford Law Review 52(5):1125–1173, URL http://www.jstor.org/stable/1229511.

Sankar, L., Rajagopalan, S., & Poor, H. (2010). Utility and privacy of data sources: Can Shannon help conceal and reveal information? *Information Theory and Applications Workshop (ITA)*, *2010*, 1–7. doi:10.1109/ITA.2010.5454092.

Scarfone, K., Mell, P. (2007). Guide to intrusion detection and prevention systems (IDPS). http://csrc.nist.gov/publications/nistir/ir7628/nistir-7628_vol1.

Schechter, S. E., & Smith, M. D. (2003). How much security is enough to stop a thief? The economics of outsider theft via computer systems and networks. *Financial Cryptography*, *2742*, 122–137.

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*(379–423), 623–656.

Smith, G. (2009). On the foundations of quantitative information flow. In: Alfaro, L. D. (Ed.), Foundations of software science and computational structures, no. 5504 in Lecture Notes in Computer Science (pp 288–302). Berlin Heidelberg: Springer.

Smith, G. (2011). Quantifying information flow using min-entropy. In: Quantitative evaluation of systems (QEST), 2011 eighth international conference on, pp 159–167, doi:10.1109/QEST.2011.31.

Sourcefire Vulnerability Research Team VRT. (2014). Download snort rules. URL http://www.snort.org/downloads/2862.

Sweeney, L. (2002). k-anonymity: A model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, *10*, 557–570.

Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. (2010). A detailed analysis of the KDD CUP 99 data set. In: Second IEEE symposium on computational intelligence for security and defence applications 2009.

Trustwave. (2014). ModSecurity open source web application firewall. URL http://www.modsecurity.org.

Ulltveit-Moe, N., Oleshchuk, V. A. (2013). Measuring privacy leakage for IDS rules. CoRR abs/1308.5421.

Warren, S., & Brandeis, L. D. (1890). The right to privacy. *Harvard Law Review*, *4*(5),

X-Force, I. S. S. (2011). Signature author's guide, IBM security systems opensignature. http://www-01.ibm.com/support/docview.wss?uid=swg21570487&aid=3.

Yannacopoulos, A. N., Lambrinoudakis, C., Gritzalis, S., Xanthopoulos, S. Z., & Katsikas, S. N. (2008). Modeling privacy insurance contracts and their utilization in risk management for ICT firms. *Proceedings of the 13th European symposium on research in computer security: Computer security* (pp. 207–222). Málaga, Spain: Springer.