

PRI: Privacy Preserving Inspection of Encrypted Network Traffic

Liron Schiff¹ Stefan Schmid²

¹ Tel Aviv University, Israel ² Aalborg University, Denmark

Abstract—Traffic inspection is a fundamental building block of many security solutions today. For example, to prevent the leakage or exfiltration of confidential insider information, as well as to block malicious traffic from entering the network, most enterprises today operate intrusion detection and prevention systems that inspect traffic. However, the state-of-the-art inspection systems do not reflect well the interests of the different involved autonomous roles. For example, employees in an enterprise, or a company outsourcing its network management to a specialized third party, may require that their traffic remains confidential, even from the system administrator. Moreover, the rules used by the intrusion detection system, or more generally the configuration of an online or offline anomaly detection engine, may be provided by a third party, e.g., a security research firm, and can hence constitute a critical business asset which should be kept confidential. Today, it is often believed that accounting for these additional requirements is impossible, as they contradict efficiency and effectiveness. We in this paper explore a novel approach, called Privacy Preserving Inspection (PRI), which provides a solution to this problem, by preserving privacy of traffic inspection and confidentiality of inspection rules and configurations, and e.g., also supports the flexible installation of additional Data Leak Prevention (DLP) rules specific to the company.

I. INTRODUCTION

The Internet has become a critical and indispensable infrastructure for many organizations. At the same time, the Internet constitutes a security threat. For example, web-based services, such as email, are indispensable for communicating with others either within or outside of an organization, but introduce the risk of data exfiltration.

Intrusion Detection Systems (IDS) as well as Intrusion Prevention Systems (IPS) are frequently used today to defend networks (or specific servers) from cyber attacks [21], [27]. In particular, such systems can prevent exfiltration of confidential insider information by blocking accidental or intentional leakage, e.g., by searching for document confidentiality watermarks in the data transferred out of an enterprise network. Such systems are also vital to control inbound traffic, and e.g., to detect if packets from a compromised sender contain an attack, employ parental filtering to prevent children from accessing adult material, etc. [25] Indeed, many cyber attacks today are carried out remotely, exploiting vulnerabilities in network components or applications, or tempting naive users to download malware and install them on their PCs.

To provide such functionality, these systems rely on traffic inspection: they allow the definition of configuration rules

which define known attack patterns, indicators for attacks, or traffic anomalies, and which are matched against the packet header *and* payload. If rules are matched, alerts are generated, and/or, in case of prevention systems or firewalls, packets are dropped. The corresponding rules can either be distributed by a local support team, by a third party (for instance a security research firm), by the network operator, or a combination thereof: third party provided rules can be complemented with organization-specific asset leakage indicators, coming from Data Leakage Prevention (DLP) systems.

While today’s intrusion detection and prevention systems typically perform well for unencrypted traffic, they struggle with encrypted traffic, resulting in false negatives or poor performance. As a workaround, in practice today, the secure and encrypted channel from or to the Internet is often terminated at a *proxy*, which essentially mounts some kind of “man-in-the-middle-attack”. While this solution ensures an effective detection and prevention, it comes at the price that the privacy of user traffic (e.g., emails) is undermined. In fact, even in the case where communication was already unencrypted anyway, performing deep packet inspection (not for the purpose of forwarding) can be seen as privacy violation. Indeed, users and clients have criticized this approach and expressed worry, e.g., that the private logged data is given to marketers [25], [29], [31].

Privacy-preserving intrusion detection may not only be desirable and relevant in the context of enterprise networks, but is also gaining in importance in the light of today’s trend to outsource the network management, including security aspects, to third parties [24]. For example, the management of third-party networks can be a lucrative business for Internet Service Providers (ISPs). At the same time, for customers running security critical businesses (for example banks), it is important that the privacy of traffic be preserved.

We in this paper however observe another confidentiality issue of today’s solutions: it concerns the confidentiality of the inspection logic itself. For example, the development and maintenance of effective intrusion detection rules is challenging, and especially small enterprises do not have the expertise and time to define the most effective rules and constantly follow the news. This constitutes a business opportunity for third parties: a company specialized into security research can take over the responsibility to define and maintain a good rule set. However, such a business model also introduces new requirements. In particular, a third party company may not be willing to share its rules,

or more generally *configurations* of (online or offline) anomaly detection systems, with the customer: these rules and configurations are an intellectual property which constitute an essential asset of the business model.

At first sight, it may appear that the requirements are contradicting: First, it seems unavoidable that an intrusion detection or prevention system, which for efficiency and effectiveness reasons needs to inspect the traffic in an unencrypted form, may leak information about the user traffic. Second, it also seems unavoidable that a system administrator operating a system based on the rules of a third party company, can see the rules.

A. Contribution

We identify the different autonomous roles involved in a traffic inspection system, including intrusion detection/prevention systems, but also more sophisticated online or offline anomaly detection systems as they may for example be required to deal with insider threats. We then explore the feasibility of providing a system which meets these requirements, respecting the autonomy of the different involved stakeholders, without introducing new threats coming, e.g., from insiders.

In particular, we present a Privacy-Preserving Intrusion detection/prevention system, short PRI (*fruit* in Hebrew), which decouples the different roles, and hence significantly reduces the required trust assumptions. PRI leverages the hardware protection of architectures like Intel SGX [30] to defend against insiders or system administrators aiming to break the confidentiality. A distinguishing feature of PRI is the simple and cheap deployment: a single trusted hardware component is sufficient. On the user side, only a simple software update is required.

B. Paper Organization

The remainder of this paper is organized as follows. In Section II, we present a model which identifies the different roles/stakeholders and the resulting requirements for the traffic inspection system. After providing the necessary background (Section III), based on these requirements, Section IV describes our proposed architecture PRI which meets these requirements. We discuss use cases in Section V and review related work in Section VI, and we conclude our contribution in Section VII. A small demo of our system is available online at <https://www.youtube.com/watch?v=b54unY8iGs0>.

II. ROLES AND REQUIREMENTS

In the following, we first present the different roles and discuss their objectives. Based on this model, we then derive security goals. Note that for ease of presentation, in the following, we will mostly focus on intrusion *detection*. However, our approach can easily be generalized to *prevention* or *configuration-based* systems, as we will elaborate more later in this paper.

We in this paper distinguish between the following roles:

- **The administrators:** The administrators are responsible for ensuring the availability and security of the network. In particular, they want to prevent the leakage of sensitive insider information, and also prevent malicious traffic entering the network from the outside. Besides relying on up-to-date security rules possibly provided by an external company, administrators may also want to be able to add their own Data Leak Prevention (DLP) rules, specific to their organization.
- **The users:** The term user will generally refer to the communication endpoints. In particular, we will usually assume that one endpoint is inside an enterprise (an *insider*) while the other is outside (an *outsider*); however, also traffic between two insiders can be subject to inspection. We assume that while users profit from a secure environment and the detection of undesired inbound and outbound traffic, they also desire a high communication performance as well as confidentiality of their traffic.
- **The rules (configuration) provider:** The rules for the intrusion detection system (or more generally the configurations for a traffic inspection system based on some open-source logic/engine) may be provided by an external security company specialized into developing and maintaining high-quality rules. The rules provider may desire that its rules remain confidential.

Given our roles and their objectives, we identify the following requirements:

- R1 **Efficient and effective inspection:** The inspection system must ensure to the administrators that relevant events and attacks will be detected successfully and fast.
- R2 **Privacy-preserving traffic inspection:** We want to ensure to the users that neither the network administrator nor the rule/configuration provider should be able to see the traffic.
- R3 **Confidentiality of rules:** We want to ensure that the rules are kept confidential and are not leaked to the other roles.

This paper is motivated by the question whether it is possible to design an architecture which meets the different goals of the different roles, maintaining their autonomy. Indeed, designing such a system seems challenging:

- Today's proxy solutions do not meet the requirements R2 and R3: a proxy server can be exploited by the administrators to learn about the unencrypted traffic. Rather, we want to develop a solution which does not allow the administrators to configure the inspection system in a way that allows them to learn details about the traffic which are not security relevant.
- Networks usually operate at very high rates: cryptographic schemes based on fully homomorphic or

functional encryption [9], [11], [13] are slow and can decrease network rates by many orders of magnitude [12], violating requirement R1.

At a first glance, the requirements seem to contradict: a system may hardly be able to efficiently and effectively inspect traffic without introducing opportunities to the administrator to see the traffic and used rules. However, as we will see in the following, there do exist solutions to satisfy these seemingly conflicting properties.

III. BACKGROUND

Before presenting our solution, we provide some background which is necessary to understand our solution. In particular, we revisit IDS systems and give an introduction to SGX.

A. Traffic Inspection Systems

Almost all cyber security breaches involve transmissions of traffic over a network. The standard approach to secure transmissions is to inspect traffic, checking whether the traffic carries the attack (e.g., a malware) or its outcome (e.g., a stolen digital asset). For this purpose, organizations deploy packet inspection systems in their networks, configuring them with known attack indicators, in the case of Intrusion Detection/Prevention Systems (IDS/IPS), and with (possibly organization-specific) asset leakage indicators in the case of Data Leakage Prevention (DLP) systems.

All inspection systems essentially search for the configured indicators inside the traffic, where indicators can be based on exact match strings, regular expressions, statistical properties, and more. Inspection of traffic in these systems usually includes the inspection of the *packet payload*, i.e., accessing the application layer data.

A simpler type of inspection that considers only the packet headers (up to the transport layer) is often performed by firewalls. A firewall is a network security system that monitors and controls the incoming and outgoing network traffic and may also include IDS/IPS capabilities. Firewalls can be considered as a barrier between a trusted, secure internal network and another outside network, such as the Internet, thereby mitigating attacks in the early stage of penetration. Firewalls can also be host based, operating on and defending a single machine.

Inspection systems are also used to detect insider threats [7], for example by analyzing confidential documents in private communications, possibly enhanced with watermarking techniques [26].

Web security today is usually realized with HTTPS, which relies on the Transport Layer Security (TLS) a.k.a. Secure Sockets Layer (SSL) protocol. TLS/SSL provides confidentiality, integrity and authentication of data in transit. The protocol offers encryption, hash functions or message digests, and digital signatures.

Encrypted traffic constitutes a great challenge to packet inspection systems as it hides the payload content from anyone but the session endpoints. There are two common solutions to this challenge. The first is to extend host based firewalls to gain access to the data *before being encrypted*: this can be implemented through integration with the host operating system. We will refer to this solution as the *client-side firewall approach*. The second solution, usually deployed by large firms, is to operate a middlebox that serves as a *proxy* between any internal user PC and any external web server. The implementation of such a solution requires to configure (trick) the user PCs to identify the proxy as the server and establish an encrypted session with it rather than with the server.

However, both the distributed firewall as well as the proxy architecture expose the user data to 3rd parties, namely the firewall producer, the proxy vendor or the network operator, possibly violating the user privacy. In addition, the inspection rules, which are the intellectual property of security researchers or firms, might be extracted at the host or the proxy, undermining their effectiveness and profit.

B. SGX

Intel *Software Guard Extensions (Intel SGX)* [8], [15], [17], [30] are new CPU instructions which allow applications to manage private regions of code and data. That is, using SGX, an application can run in a protected environment, the so-called *enclave*, secure from malware or the inspection by the computer administrators.

There is no need to encrypt the protected portion of an application for distribution. Before the enclave is built, the enclave code and data is free for inspection and analysis. When the protected portion is loaded into an enclave, its code and data is measured. An application can prove its identity to a remote party and be securely provisioned with keys and credentials. The application can also request an enclave and platform specific key that it can use to protect keys and data that it wishes to store outside the enclave. In addition to the security properties, the enclave environment offers scalability and performance associated with the execution on the main CPU of an open platform.

IV. PRIVACY PRESERVING INSPECTION

This section shows that the design goals and requirements derived above can actually be met, assuming the availability of a trusted hardware. In particular, this section presents our Privacy Preserving Inspection architecture, short PRI. PRI allows the inspection, and if considered harmful, possibly prevention, of encrypted traffic while guaranteeing that no information about the traffic is leaked from the inspecting device, and that inspection rules are not revealed to the operator of the device nor the traffic generator.

In a nutshell, the idea underlying PRI is to decouple and separate the different roles (e.g., users, administrators,

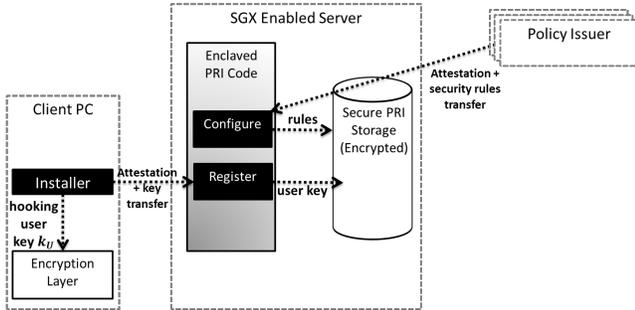


Figure 1. Setup of the PRI system components. The PRI process is configured to inspect user traffic according to security rules inserted by the policy issuer. The process is secured and ensures confidentiality of traffic and rules. Keys and rules are stored safely in the PRI storage.

security company) by defining an interface between them and supporting verification. In order to achieve this decoupling, PRI relies on a single device (such as Intel SGX or ARM TEE) which decrypts traffic from the users, applies the desired security rules, and when needed raises alerts or drops the traffic, without leaking information about the user traffic nor the applied rules, to the roles which are not allowed to have this information. Interestingly, a single trusted box is sufficient for PRI: no hardware modifications are required on the user side. Moreover, an important aspect of PRI is that the user side agent and the enclaved system can be *open-source*: thus, the users and the research community can and should verify its privacy preserving property. The attestation allows the agent to verify that it communicating with the verified system code.

A. Setting Up PRI

Let us first discuss the setup of PRI (cf Figure 1). We describe it for the enterprise intrusion detection scenario, but it is easy to adapt it to ISP outsourcing scenarios. The PRI system consists of the user’s communication device, e.g., a personal computer (PC) (or smartphone, laptop, etc.) and the PRI server which need to be set up as follows. On the user side, we simply need to install a PRI agent at the encryption layer of the OS or specific application (e.g., web browser). The PRI agent is configured with a user key, k_U , which is securely transmitted (through attestation) to the enclaved PRI system. In addition the agent is installed in a way that allow it to gain the key, k_S , of every encrypted session S , that uses the encryption layer enabling it to securely send it to the PRI system.

The setup also involves configuring the PRI server with the rules (also called the policy) that need to be checked against the traffic. Rules can be configured by multiple policy issuers, each of them can connect securely to the PRI process, validating its secure execution (attestation) and sending the rules. As a result, the PRI process stores the rules in the secured storage, to be used when a user traffic needs to be inspected.

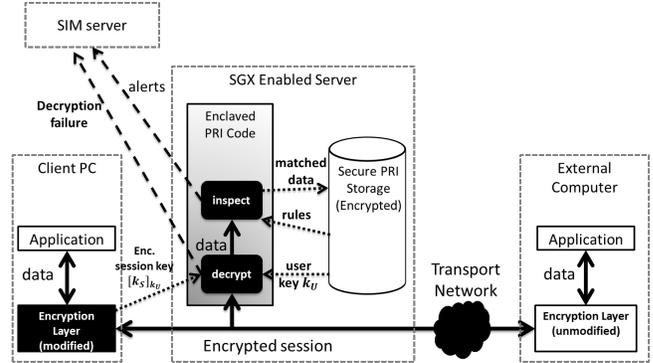


Figure 2. The PRI system inspects user session using the key sent by the agent. Matches are saved in a secured storage and alerts are sent to the corporate SIM server about detected anomalies.

B. Operating PRI

Once the system has been set up, the registered users’ traffic can be inspected. Inspection is carried in the following steps, see also Figure 2.:

- 1) An encrypted session is established between the applications of the user PC and an external (unregistered) or internal PC.
- 2) The user PC connects to the PRI process (at the PRI server), and sends to it the session key encrypted by the shared user key, i.e., $[k_S]_{k_U}$.
- 3) The session traffic is duplicated and processed by the PRI process.
- 4) The session traffic is decrypted, using the session key.
- 5) The session data (clear text) is inspected with the rules configured in the system.
- 6) In case part of the data matches a rule, this part is securely stored in the PRI storage and an alert (containing the rule identifier) is reported the corporate Security Information Management (SIM) server.

Note that in case of misconfiguration, or malicious activity in the user PC that results in decryption failure at the PRI system, a special alert is reported.

There exists another more subtle threat for the confidentiality of the traffic: even if the traffic cannot be inspected, the reported matches over traffic leak information. In order to prevent an attacker from abusing the security rules to learn about private user data (e.g., matching every possible byte or word), the PRI system allows each user to request the parts of his or her traffic which have been matched by the inspecting device, and thus to learn about and detect the use of abnormal matching rules. To ensure this, our system uses a special viewer application which connects to the PRI process, authenticates as the user and issues the request. As a result, the PRI process retrieves the user’s matches from the secured storage and sends them to the viewer app (cf Figure 3.).

In addition, the PRI system can securely inspect static (e.g., likelihood of match for common words) and dynamic (e.g.,

		Use Case		
		Enterprise Security and Insiders	Outsourced Security	Anti-Terror Intelligence
Entities	Clients	employees	enterprise	civilians / web hosts
	PRI Operators	admins	external	ISPs
	Rule Providers	admins + external	external	governments

Table I
COMPARISON OF PRI USE CASES.

today. In particular, such services can be used by a malicious insider to steal intellectual property or sensitive company information. Indeed, insider attacks pose some unique challenges for security administrators. In general, it can be challenging to configure an intrusion detection system and define a good rule set to detect internal attacks: different users should have access to and use many different services, servers, and systems for their work. However, over the last years, several interesting IDS-based solutions have been developed to detect insider threats [7], e.g., searching for confidential documents and watermarks in private communications [26]. Our approach is directly applicable to these systems.

Even in scenarios where IDS-based systems may be insufficient, PRI can be attractive: PRI is also applicable in more general security solutions, for example systems which are based on machine learning or anomaly detection of online or offline traffic data, and which can consist of an open-source engine and a possibly confidential configuration. More concretely, advanced methods to detect insider attacks usually consider a wider context when processing network events, therefore requiring to store and query the entire event history [16], and include machine learning algorithms [5] (e.g., anomaly detection). Our design can be extended to support these methods, by utilizing the secure storage to save the events history and to execute the advanced inspection as an enclaved process that analyzes the history instead of the current traffic. In these cases, our extended system also allows to split the open source inspection engine from the inspection rules, e.g., the big data queries or the machine learning parameters and filters, while supporting very general notions of privacy [20], [18], [19].

B. Network Outsourcing

Another use case arises in the context of network outsourcing: We currently witness the trend that enterprises wish to outsource their cyber security logic or even the entire network administration to an external company experienced in this field. Translating this scenario to our framework, the enterprise (or any of its employees and their PCs) constitutes the client, executing the PRI agent, and the external company operates the PRI server, executing the enclaved (and open-source) PRI software.

C. Anti-Terror Intelligence

An interesting use case also arises in the context of governments gathering intelligence from their own citizens

in order to fight terror. This use case is highly controversial today and subject to major ongoing debates [1], due to the tradeoff of preserving civil liberties (i.e., the right for privacy) and the government duty to save lives. With the PRI system, governments and citizens can decide and control the level of inspection and privacy infringement applied to the private data. The anti-terror intelligence use case may require some adaptations to the model. For example, the PRI system can be operated by the ISPs that (as part of the PRI software) can correlate the traffic with client information. The rules are securely provided by the government, but some aggregated information on them can be made accessible to the public, or its representatives. The exact matches of user traffic, cannot be provided immediately to the user as this might prevent the authorities to effectively react to threats, therefore the PRI system may make use of a notification delay which is defined a priori.

VI. DISCUSSION AND RELATED WORK

In an age where more and more resources and services are outsourced, there is an increasing need for solutions preserving critical security aspects. Over the last years, this problem has been discussed particularly intensively in the context of cloud computing: Although users of cloud computing infrastructure may expect their data to remain confidential, today's clouds are built using a classical hierarchical security model that aims only to protect the privileged code (of the cloud provider) from untrusted code (the user's virtual machine), and does nothing to protect user data from access by privileged code [3]. While for many large-scale computations today, the use of cloud computing resources is unavoidable or at least financially very attractive, users may not be willing to trust their cloud provider to keep their data confidential. In fact, the cloud user must trust not only the hardware on which her or his data is actually analyzed, but in addition also (i) the provider's software, including privileged software such as a hypervisor and firmware but also the provider's full stack of management software; (ii) the provider's staff, including system administrators but also those with physical access to hardware such as cleaners and security guards; (iii) the law enforcement bodies in any jurisdiction where their data may be replicated, as the Snowden leaks have revealed [3], [10].

Recently, researchers have started investigating whether similar approaches as proposed for cloud computing may also be applicable in the context of computer networks, which

		Architecture			
		Proxy	Client-Side Firewall	BlindBox [25]	PRI
Privacy	user exposure	middlebox	no	no	no (enclaved)
	rules exposure	middlebox	endpoint	middlebox	no (enclaved)
Effectiveness	inspection guarantee	yes	no	no	yes
	supports rules	any	any	exact match only	any
Overhead	computation	en-&decryption	none	tokenization & encryption	decryption
	communication	none	none	a stream of tokens	one packet

Table II
COMPARISON OF PERFORMANCE AND SECURITY OF DIFFERENT IDS ARCHITECTURES.

come with rather different requirements. Existing systems are based on proxies and are vulnerable to a man-in-the-middle attack on SSL, installing fake certificates at the middlebox [14], [23]. The middlebox can break the security of SSL and decrypt the traffic so it can perform the Deep-Packet Inspection (DPI). The removal of the SSL end-to-end security, results in a host of issues. Some proposals allow users to tunnel their traffic to a third party middlebox provider [22], [24]. But these approaches allow the middlebox owner to inspect/read all traffic. An alternative today are distributed firewalls, a client-side approach to implement intrusion detection/prevention.

A very interesting approach is taken by BlindBox [25]: unlike PRI, BlindBox performs the deep-packet inspection directly on the *encrypted traffic*. In a nutshell, in BlindBox, the endpoint generates a tokenized version of the traffic which can be inspected in privacy preserving manner at a special server. This tokenized traffic is sent in parallel to the origin, increasing the load in the network. However, BlindBox requires the user to compute hashes of traffic segments and to send them to the inspection box, thereby introducing computational overhead to the user PC, as well as traffic overhead to the network. Moreover it is dependent on the cooperation of the user PC to perform its part of the scheme. In addition, BlindBox only supports exact match rules and not regular expressions, that are commonly used in security policies.

We believe that our approach nicely complements these works, and focuses on a relevant use case, namely traffic inspection. An SGX approach as recently suggested in the context of cloud computing, could be used to implement our secure server, and the decoupling of the various roles identified in our paper.

Table II summarizes the advantages and disadvantages of the different architectures: the man-in-the-middle proxy and the client-side distributed firewalls used today, as well as BlindBox and PRI. Among these solutions, only BlindBox and PRI are providing the required privacy guarantees. The main limitation of BlindBox is arguably its expressive power, while the main limitation of PRI is its dependency on a SGX hardware (although a single box is sufficient).

VII. CONCLUSION

This paper studied the classic problem of traffic inspection from an interesting new, privacy preserving perspective. In particular, while today it is commonly believed that it is inevitable that users have to blindly trust the administrator managing the intrusion detection or prevention system, we in this paper, have questioned this assumption. In particular, we have shown that it actually is possible to reduce trust assumptions in the enterprise network, and presented an intrusion detection system which is not only privacy preserving regarding the user traffic but also regarding the rules used in the IDS/IPS. Interestingly, the proposed PRI system requires a single secured server; no modifications of the hardware at the users is required.

In summary, the PRI system features the following properties:

- 1) It decrypts and inspects network traffic in a privacy preserving manner.
- 2) It accepts new security rules from administrators and applies them to the traffic in a secure and privacy-preserving manner: the inspecting device does not leak any unnecessary information about the user traffic.
- 3) It can be configured with new rules from rule providers in a secure and privacy-preserving manner: the inspecting device does not leak any unnecessary information about the user traffic.

We believe that the ideas underlying PRI can be used in several contexts and beyond the use cases discussed in this paper. In particular, it is not limited to rule-based intrusion detection systems, but can also be useful in the context of more sophisticated and offline systems, as they may for example be required to handle advanced insider threats [2], [6], [7].

Our approach raises several interesting questions for future research. Obviously, the performance of our architecture needs to be evaluated in detail. Moreover, it will be interesting to explore further the applications of trusted execution environments in the context of computer networking and network function virtualization.

We have started implementing a simple test version of the PRI intrusion detection/prevention system. Due to the delayed introduction of SGX hardware, the code is still very simple and contains untrusted parts.

REFERENCES

- [1] H. Abelson, R. Anderson, S. M. Bellovin, J. Benaloh, M. Blaze, W. W. Diffie, J. Gilmore, M. Green, S. Landau, P. G. Neumann, R. L. Rivest, J. I. Schiller, B. Schneier, M. A. Specter, and D. J. Weitzner. Keys under doormats. *Commun. ACM*, 58(10):24–26, Sept. 2015.
- [2] R. Barrios. A multi-leveled approach to intrusion detection and the insider threat. *Scientific Research Publishing*, 2013.
- [3] A. Baumann, M. Peinado, and G. Hunt. Shielding applications from an untrusted cloud with haven. *ACM Trans. Comput. Syst.*, 33(3), Aug. 2015.
- [4] M. Belshe, R. Peon, and M. Thomson. Hypertext transfer protocol version 2 (http/2). RFC 7540, RFC Editor, May 2015. <http://www.rfc-editor.org/rfc/rfc7540.txt>.
- [5] O. Brdiczka, J. Liu, B. Price, J. Shen, A. Patil, R. Chow, E. Bart, and N. Ducheneaut. Proactive insider threat detection through graph learning and psychological context. In *Security and Privacy Workshops (SPW), 2012 IEEE Symposium on*, pages 142–149, May 2012.
- [6] D. Cappelli, A. Moore, R. Trzeciak, and T. Shimeall. Common sense guide to prevention and detection of insider threat. *CERT Insider Threat Study Team, Carnegie Mellon University*, 2008.
- [7] N. Einwechter. Preventing and detecting insider attacks using ids. *SecurityFocus, March*, 2002.
- [8] F. McKeen. Innovative instructions and software model for isolated execution. In *Proc. HASP*, 2013.
- [9] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *Proc. IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 40–49, 2013.
- [10] B. Gellman and A. Soltani. Nsa infiltrates links to yahoo, google data centers worldwide, snowden documents say. In *The Washington Post*, 2013.
- [11] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 169–178, 2009.
- [12] C. Gentry, S. Halevi, and N. P. Smart. Homomorphic evaluation of the AES circuit. In *Proc. 32nd Annual Cryptology Conference Advances in Cryptology (CRYPTO)*, pages 850–867, 2012.
- [13] S. Goldwasser, Y. Kalai, R. A. Popa, V. Vaikuntanathan, and N. Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Proc. 45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 555–564, 2013.
- [14] L. S. Huang, A. Rice, E. Ellingsen, and C. Jackson. Analyzing forged ssl certificates in the wild. In *Proc. IEEE Symposium on Security and Privacy (SP)*, pages 83–97, 2014.
- [15] I. Anati et al. Innovative technology for cpu based attestation and sealing. In *Proc. HASP*, 2013.
- [16] S. Inc. Detecting insider threats: How splunk software is used to safeguard financial data, 2014.
- [17] M. Hoekstra et al. Using innovative instructions to create trustworthy software solutions. In *Proc. HASP*, 2013.
- [18] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian. L-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1), Mar. 2007.
- [19] F. D. McSherry. Privacy integrated queries: An extensible platform for privacy-preserving data analysis. In *Proc. ACM SIGMOD International Conference on Management of Data (SIGMOD)*, pages 19–30, 2009.
- [20] A. Meyerson and R. Williams. On the complexity of optimal k-anonymity. In *Proc. 23rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 223–228, 2004.
- [21] V. Paxson. Bro: A system for detecting network intruders in real-time. *Comput. Netw.*, 31(23-24), Dec. 1999.
- [22] A. Rao, J. Sherry, A. Legout, A. Krishnamurthy, W. Dabbous, and D. Choffnes. Meddle: Middleboxes for increased transparency and control of mobile traffic. In *Proc. ACM CoNEXT*, pages 65–66, 2012.
- [23] Runa. Security vulnerability found in cyberoam dpi devices (cve-2012-3372). In *Tor Project Blog*, 2012.
- [24] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. Making middleboxes someone else’s problem: Network processing as a cloud service. In *Proc. ACM SIGCOMM*, pages 13–24, 2012.
- [25] J. Sherry, C. Lan, R. A. Popa, and S. Ratnasamy. Blindbox: Deep packet inspection over encrypted traffic. In *Proc. ACM SIGCOMM*, pages 213–226, 2015.
- [26] G. Silowash, T. Lewellen, J. Burns, and D. Costa. Detecting and preventing data exfiltration through encrypted web sessions via traffic inspection. Technical Report CMU/SEI-2013-TN-012, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, 2013.
- [27] Snort. website. <https://www.snort.org/>, 2016.
- [28] Suricata. website. <http://suricata-ids.org/>, 2016.
- [29] N. Vallina-Rodriguez, S. Sundaresan, C. Kreibich, N. Weaver, and V. Paxson. Beyond the radio: Illuminating the higher layers of mobile networks. In *Proc. ACM MobiSys*, pages 375–387, 2015.
- [30] S. Website. <https://www.sans.org/reading-room/whitepapers/authentication/ssl-tls-whats-hood-34297>. 2016.
- [31] K. Zetter. The feds cut a deal with in-flight wi-fi providers, and privacy groups are worried. In *Wired Magazine*, 2014.