Reactive and Adaptive Security Monitoring in Cloud Computing

Clément Elbaz, Louis Rilling, Christine Morin

► To cite this version:

Clément Elbaz, Louis Rilling, Christine Morin. Reactive and Adaptive Security Monitoring in Cloud Computing. FAS* Doctoral Symposium 2018, Sep 2018, Trento, Italy. pp.1-3, 10.1109/FAS-W.2018.00014 . hal-01884739

HAL Id: hal-01884739

<https://hal.inria.fr/hal-01884739>

Submitted on 1 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Reactive and Adaptive Security Monitoring in Cloud Computing

Clément Elbaz
Univ Rennes, Inria, CNRS, IRISA
Rennes, France
clement.elbaz@inria.fr

Louis Rilling
DGA
Rennes, France
louis.rilling@irisa.fr

Christine Morin
Univ Rennes, Inria, CNRS, IRISA
Rennes, France
christine.morin@inria.fr

***Index Terms*—Security, Cloud, SLA, IDS**

I. MOTIVATION

Cloud computing enabled service-level agreements (SLAs) to gain widespread use among information systems stakeholders. It is now normal for performance and availability of such systems to be carefully measured and evaluated. Contracts that include financial penalties in case of breach are now common.

However security is lagging behind this trend; it is as important to stakeholders as performance and availability, but is generally not included in the scope of service-level agreements between stakeholders, and handled instead on a best-effort basis, without any transparency nor SLA with their clients.

One reason for this is the difficulty of objectively measuring security. Indeed, the actual security level of a system is dependent on a wide range of factors, some intrinsic to the system - such as a design or implementation mistake resulting in a vulnerability - and some extrinsic to it. For instance, an external event such as the publication of a vulnerability in an open-source software dependency or a change of political context in a country can widely impact the risks faced by an information system even if no actual change were made to the system.

These factors are even more numerous in multi-tenant cloud infrastructures because of the sheer number of actors involved - and their sometimes conflicting incentives - and opportunities for both attack and defense at scale.

Security monitoring aims to detect and react to attacks in real time; Reactive security monitoring intends to take external events into account while doing so. Improving the accuracy of a system's overall security assessment can help transitioning security to a SLA paradigm and enable better transparency for all stakeholders.

II. OBJECTIVES

We present an intermediate stage research project whose goal is threefold. First we aim to better define the various threats looming over information systems in the cloud, more specifically the way these threats evolve with time. We are particularly interested in the software vulnerability life cycle and how newly found vulnerabilities affect the security of cloud tenants. Similarly, we want to study how cloud providers

can use intrusion detection systems (IDS) - which detect the use of a vulnerability through network analysis - and corrective patches - which fix the actual vulnerability through a software upgrade - to mitigate the threat induced by these vulnerabilities to tenant information systems; and provide this mitigation in a controlled time frame.

Second we want to provide new types of high level security SLAs allowing cloud providers to make these adaptive changes to their infrastructure in a transparent and observable way for their clients.

Third we are interested in the economic incentives of stakeholders regarding reactive security. We want to get a better understanding of the ways cloud customers and providers spend (or don't spend) their resources on security. We are also interested in studying how various business models can encourage (or discourage) adoption of responsible security practices among cloud stakeholders.

III. METHODOLOGY

In order to accomplish the goals described in Section II, we need to overcome two challenges.

First, there is a lack of consensus on the best way to evaluate the actual security of an information system [1] [2], and more specifically evaluating the efficiency of IDSs [3]. Indeed, evaluation datasets such as the DARPA dataset [4], NSL-KDD [5] and UNIBS datasets [6] are questioned on their realism as well as their currentness. Meanwhile evaluations based on actual network traffic have trouble separating normal and malicious activities and are difficult to replicate after the fact.

A second challenge is the difficulty of working with new threats: reactive security monitoring must deal with ill-defined situations and makes decisions while most of the information is not yet available.

Our answer to these challenges is to consider reactive security as an incomplete information game exhibiting some uncertainty.

Therefore we aim to approach reactive security through game theory and economic modeling. By acknowledging the uncertainty in our models, we expect to get a more accurate picture of the expected value of a security situation as well as a better understanding of why (and if) we should be confident about our analysis.

Security viewed through the prism of game theory and economics has already received attention [7] [8] and we plan to build our work upon these publications. Likewise, the existing work related to the software vulnerability lifecycle [9] [10] [11] [12] is also a foundation of our own.

IV. RESEARCH PLAN

Our first work focused on a specific aspect of reactive security: studying the way an IaaS cloud service provider (CSP) can react to newly published vulnerabilities. More specifically we studied the possibility for a CSP to engage in a contract with its clients supported by a new type of SLA. This SLA covers the ability for the CSP to provide counter-measures for any new vulnerability affecting the client, and do so in a contractual time frame. The contract includes a periodic (monthly for instance) payment from the client to the CSP, and a financial penalty from the CSP to the client in case of a breach of SLA.

We focused on two types of counter-measures: corrective patches deployed in a software repository (such as the debian-security repository from the Debian Linux distribution [13]) and IDS signatures rules (such as the official Snort rules written by Talos [14]).

This led us to carry out a new vulnerability life cycle study, the first of its kind encompassing IDS signatures rules as well as correctives patches. For this study we oversaw the lifecycle of 34 000 vulnerabilities published between June 2014 and October 2017.

Then we used the data from this study to analyze the business model of this contract.

This study led to promising results: we showed that for a wide range of plausible scenarios, it was possible for a CSP to make a profit from charging its clients to enforce a SLA on counter measure deployment. Those results were published in July 2018 [15].

However this study also showed the lack of suitability of rule-based IDSs for reactive security. Indeed, rule-based IDSs use static, human authored signatures to detect attacks. In our study we noticed that most of these rules were published long after the associated corrective patches. This led us to conclude that signature-based IDSs are generally inadequate when dealing with very recent threats.

Therefore we chose anomaly-based IDSs as our next research topic. Anomaly-based IDSs are machine-learning based systems attempting to separate normal and abnormal network activity. Unlike signature-based IDSs they require a training period to learn what the normal state of the system is supposed to be. Our hypothesis is that anomaly-based IDSs might be better suited than signature-based IDSs when dealing with new vulnerabilities, as they don't require manual rule authorship. We plan to spend most of 2018 evaluating them theoretically and experimentally to assess their suitability to reactive security. As mentioned in Section III, one challenge is the lack of a consensus on the best way to evaluate IDSs. Another challenge more specific to anomaly-based IDSs is the lack of open source software widely used in production: most

of anomaly-based IDSs are prototypes created for specific research projects, making their real world evaluation difficult. This situation stems in part from the high false positive rate of current anomaly-based IDS designs. Nevertheless we plan on experimenting with Hogzilla [16] and possibly our own prototype.

Another topic we plan to work on is to provide a new high level security SLA that can be translated into evolving low level SLAs. In particular, we view SLAs as a bridge between the technical and organizational aspects of security. A high level SLA can define contractual obligations for both the CSP and the client, as well as defining high level technical objectives that can be translated into context-dependent low level requirements. We plan to study the state of the art on this topic and possibly propose new paradigms.

In 2019 we plan to gather what we learned in 2017 and 2018 to design and implement a prototype of a solution capable of making reactive decisions to external threats. This prototype will combine the use of both corrective patches and signature and anomaly-based IDSs to exploit the strength of both approaches. We will experimentally evaluate this prototype on the Grid5000 platform [17] for both performance and security aspects.

REFERENCES

- [1] b. K. Bernsmed and M. G. Jaatun and P. H. Meland and A. Undheim, "Security SLAs for Federated Cloud Services," Aug 2011, pp. 202–209.
- [2] A. Dulaunoy. (2016) The Myth of Software and Hardware Vulnerability Management. https://www.foo.be/2016/05/The_Myth_of_Vulnerability_Management/. [Online]. Available: https://www.foo.be/2016/05/The_Myth_of_Vulnerability_Management/
- [3] M. H. Bhuyan and D. K. Bhattacharyya and J. K. Kalita, "Network Anomaly Detection: Methods, Systems and Tools," *IEEE Communications Surveys Tutorials*, vol. 16, no. 1, pp. 303–336, First 2014.
- [4] DARPA Intrusion Detection Data Sets. <https://www.ll.mit.edu/ideval/data/>. [Online]. Available: <https://www.ll.mit.edu/ideval/data/>
- [5] NSL-KDD dataset. <http://www.unb.ca/cic/datasets/nsll.html>. [Online]. Available: <http://www.unb.ca/cic/datasets/nsll.html>
- [6] University of Brescia dataset. <http://netweb.ing.unibs.it/~ntw/tools/traces/>. [Online]. Available: <http://netweb.ing.unibs.it/~ntw/tools/traces/>
- [7] A. Barth, B. I. P. Rubinstein, M. Sundararajan, J. C. Mitchell, D. X. Song, and P. L. Bartlett, "A Learning-Based Approach to Reactive Security," *CoRR*, vol. abs/0912.1155, 2009. [Online]. Available: <http://arxiv.org/abs/0912.1155>
- [8] S. Zhang, X. Zhang, and X. Ou, "After We Knew It: Empirical Study and Modeling of Cost-effectiveness of Exploiting Prevalent Known Vulnerabilities Across IaaS Cloud," in *Proceedings of the 9th ACM Symposium on Information, Computer and Communications Security*, ser. ASIA CCS '14. New York, NY, USA: ACM, 2014, pp. 317–328. [Online]. Available: <http://doi.acm.org/10.1145/2590296.2590300>
- [9] S. Frei, M. May, U. Fiedler, and B. Plattner, "Large-scale Vulnerability Analysis," in *Proceedings of the 2006 SIGCOMM Workshop on Large-scale Attack Defense*. New York, NY, USA: ACM, 2006, pp. 131–138. [Online]. Available: <http://doi.acm.org/10.1145/1162666.1162671>
- [10] S. Clark, S. Frei, M. Blaze, and J. Smith, "Familiarity Breeds Contempt: The Honeymoon Effect and the Role of Legacy Code in Zero-day Vulnerabilities," in *Proceedings of the 26th Annual Computer Security Applications Conference*. New York, NY, USA: ACM, 2010, pp. 251–260. [Online]. Available: <http://doi.acm.org/10.1145/1920261.1920299>
- [11] M. Shahzad, M. Z. Shafiq, and A. X. Liu, "A Large Scale Exploratory Analysis of Software Vulnerability Life Cycles," in *Proceedings of the 34th International Conference on Software Engineering*. Piscataway, NJ, USA: IEEE Press, 2012, pp. 771–781. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2337223.2337314>

- [12] L. Bilge and T. Dumitras, “Before We Knew It: An Empirical Study of Zero-day Attacks in the Real World,” in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*. New York, NY, USA: ACM, 2012, pp. 833–844. [Online]. Available: <http://doi.acm.org/10.1145/2382196.2382284>
- [13] Debian – Security Information. <https://www.debian.org/security/>. [Online]. Available: <https://www.debian.org/security/>
- [14] Talos - Author of the Official Snort Rule Sets. <https://snort.org/talos>. [Online]. Available: <https://snort.org/talos>
- [15] C. Elbaz, L. Rilling, and C. Morin, “Mesurer et prévenir l’évolution de la menace dans un cloud d’infrastructure,” in *ComPas’18*, Toulouse, France, Jul. 2018. [Online]. Available: <https://hal.inria.fr/hal-01816674>
- [16] Hogzilla IDS. <http://ids-hogzilla.org/>. [Online]. Available: <http://ids-hogzilla.org/>
- [17] D. Balouek, A. Carpen Amarie, G. Charrier, F. Desprez, E. Jeannot, E. Jeanvoine, A. Lèbre, D. Margery, N. Niclausse, L. Nussbaum, O. Richard, C. Pérez, F. Quesnel, C. Rohr, and L. Sarzyniec, “Adding Virtualization Capabilities to the Grid’5000 Testbed.”