Relationships Between Information Security Metrics: An Empirical Study

Rodrigo Sanches Miani
School of Electrical and Computer Engineering
University of Campinas
Campinas, SP, Brazil
rsmiani@decom.fee.unicamp.br

Michel Cukier
A. James Clark School of Engineering
University of Maryland
College Park, MD, USA
mcukier@umd.edu

Bruno Bogaz Zarpelão
School of Electrical and Computer Engineering
University of Campinas
Campinas, SP, Brazil
bzarpe@decom.fee.unicamp.br

Leonardo de Souza Mendes
School of Electrical and Computer Engineering
University of Campinas
Campinas, SP, Brazil
lmendes@decom.fee.unicamp.br

ABSTRACT

Finding relevant metrics in information security is an important but difficult problem. In this paper, we propose to empirically investigate the relevance of different security metrics that could be derived from intrusion prevention system (IPS) alert events and computer security incident data. Based on the data provided by the University of Maryland, we show that IPS metrics are linked to security incidents, and also that different types of security incidents have different significant metrics. These results can be used for identifying possible candidates for security incident indicators, developing methods to improve incident prevention and helping organizations interpret their IPS's better in the future.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General—*Security and protection (e.g., firewalls)*; K.6.5 [**Management of Computing and Information Systems**]: Security and Protection

General Terms

Measurement, Experimentation, Security.

Keywords

Network and Security Management, Security Metrics, Empirical Study, Security Incidents, Intrusion Prevention Systems.

1. INTRODUCTION

In information security, questions such as "Is security improving over time?" and "Are we using effective controls?" could be used to derive measurements to facilitate decision making and improve performance and accountability. There are many suggestions in the security community for what measures organizations should collect to construct security metrics models [8], [7], [6], [12]. However, as pointed out by Jansen [5] much of what has been written about security quantification is definitional, aimed at providing guidelines for defining a security metric and specifying criteria that need to be fulfilled. According to Verendel [9] for most cases it is unknown whether the proposed models are valid or not in representing security for systems in realistic environments due to the lack of validation. In other words, research is needed to validate connections between measures and security, and determine possible correlations.

The motivation of this paper is to provide some insight to help security analysts extract relevant information about the organization security by investigating which metrics are most indicative of the occurrence of a security incident. We based our study on an empirical analysis of relationships between information security metrics. We study connections between metrics derived from intrusion prevention systems (IPS) alert events and the number of security incidents. An intrusion prevention system is an extension of an intrusion detection system (IDS) that monitors malicious activity and reacts in real time by blocking a potential attack. Such systems can be also seen as monitors of malicious activity inside and outside the organization. According to the US-CERT (United States Computer Emergency Readiness Team), an incident is the act of violating an explicit or implied security policy. Usually, every successful incident (or attack) reported is recorded in a security incidents database. In other words, if successful attacks could be found in one database and malicious activities in another database, the question arises: are malicious activities related with successful attacks? If so, this information can be used to improve the comprehension of organization security, and also to derive new attack prevention rules.

First of all, we define two sets of IPS metrics according to the attacker's perspective: i) where a computer outside the organization is targeting the organization and ii) where a computer inside the organization is targeting computers outside the organization. Since incidents greatly differ, we propose to group them into three groups: Hacking, Bot and Spam. We then empirically examine, using a multiple linear regression model, the effects of IPS metrics on the number of security incidents reported by the Division of Information Technology at the University of Maryland.

This paper is structured as follows. Section 2 describes the re-

CSIIRW '12, Oct 30 - Nov 01 2012, Oak Ridge, TN, USA
ACM 978-1-4503-1687-3/12/10.

lated work. Section 3 introduces the empirical modeling approach and presents the regression results. Section 4 provides the summary of results and discusses the impact of exploring connections between IPS metrics and security incidents. Section 5 discusses the threats to validity of our study. We provide conclusions and directions for future work in Section 6.

2. RELATED WORK

Wang and Wulf [10] introduced a framework for measuring system security based on a selection of units and scales, definition of an estimation methodology and validation of measures. Researchers have also proposed security measurement models based on attack graphs [11], attack surface [4] and risk assessment [1]. Jansen [5] provides an overview of security metrics area and discusses possible research avenues. The author states that much of what has been written about security metrics is definitional, aimed at providing guidelines for defining a security metric and specifying criteria to fulfill. A list of security metrics and data definitions that can be used across organizations to collect and analyze data on security process performance and outcomes can found in [7], [6] and [2]. However, little has been reported on actual metrics that have been proven useful in practice. Chrun et al. [3] presents a method for retrieving useful information from imperfect IPS event data using security metrics. The authors introduced an approach that consists in analyzing the evolution of IPS metrics per attack type group by focusing on outliers. Our work focus on extracting empirical relationships between IPSs and computer security incidents datasets using an approach based on security metrics. We also investigate how to use these results to improve the knowledge about system security.

3. EMPIRICAL MODELING APPROACH

The proposed IPS metrics are based on Chrun et al. [3] work, according to two security perspectives: 1) where a computer outside the organization is targeting the organization, and 2) where a computer inside the organization is targeting computers outside the organization. In both cases, we want to assess the volume and the nature of the malicious activity.

We introduce the following six IPS metrics for attacks towards the organization, named G1 metrics: i) *Number of alerts:* illustrates the number of attack attempts towards the organization, ii) *Number of (distinct) attackers:* reflects the activity of attacks targeting the organization, iii) *Number of (distinct) targets*: indicates the number of targeted computers and thus potentially future corrupted computers, iv) *Number of (distinct) signatures*: reveals the spectrum of attacks that target the organization, v) *Number of alerts per target:* indicates how much some computers are being targeted and vi) *Number of attackers per target:* reveals some possible coordinated attacks from many attackers against specific targets. When focusing on the traffic originating inside the organization and targeting computers outside the organization, we define the following six IPS metrics, named G2 metrics: i) *Number of alerts:* indicates the number of attack attempts from inside the organization to launch attacks against targets outside the organization, ii) *Number of (distinct) attackers:* illustrates the activity of attacks originating inside the organization and targeting computers outside the organization, possibly reflecting the number of corrupted computers inside the organization, iii) *Number of (distinct) targets:* reflects the number of targeted computers from attacks originating inside the organization, iv) *Number of (distinct) signatures:* illustrates the range of attack types towards targets outside the organization, v) *Number of alerts per attacker:* reflects how much computers of the organization are attacking computers outside the organization and vi) *Number of targets per attacker:* helps characterizing the malicious activity originating inside the organization.

Table 1: Summary of incident data

Incident category	Start date	End date	# of Incidents	# of Days
All	01/02/2007	12/31/2010	1776	581
Bot	01/02/2007	12/14/2010	332	144
Hacking	01/02/2007	12/31/2010	665	308
Spam	01/04/2007	12/31/2010	663	233

Table 2: Summary of IPS data

IPS group	# of Alerts	# of Attackers	# of Targets	# of Signatures
All	7,665,085	1,837,321	368,100	24,354
G1	7,293,795	1,815,492	304,068	18,265
G2	371,290	21,829	64,032	6089

According to the Center of Internet Security [8] it is possible to extract several security incident metrics from a security incident database, such as: mean time to incident discovery, mean time between security incidents and number of security incidents. In this work, we will focus on the study of the number of security incidents.

3.1 Data and Measurements

The dataset provided by the University of Maryland consists of over 1794 security incidents and 7,665,085 IPS alerts recorded during a period of four years (from January 1, 2007 to December 31, 2010). The data were grouped in a weekly basis ($t = 209$ weeks), with Monday as the first day of the week. Grouping the data in a time window is a technique to minimize the effects of lag time between occurrence of an event and submission of an incident report. However, due to the human interaction in the incident reporting process, we cannot prove that an incident reported in certain week would have been related to IPS alerts from the same week.

The incidents recorded were based on three sources of events: 1) an IDS, 2) reports from users and 3) reports from other system administrators. Since recorded incidents led to the blocking of the suspected computer's IP address, the Division of Information Technology (OIT) verified the authenticity of each incident. Therefore, all incidents obtained from these three sources were manually reviewed. OIT launched port scans and packet captures to validate the suspicious behavior of identified hosts. Based on the IDS rule that raised an alert, about 60% of the alerts were inconclusive. Among the remaining 40%, about half led to the direct action of OIT blocking the IP address and half required a confirmation. Among the incident alerts, very few were reports from users. The reports from other system administrators were defined as incidents in roughly 75% of the cases based on the source trustworthiness.

The incidents dataset includes seven different categories: Abuse, Bot, Hacking, Spam, Virus, Spyware and Worm. Each category contains several incident types, involving a total of 52 types. We examined the three incident types that occurred most frequently in the dataset: Bot, Hacking and Spam. These three incident types account for 1660 of a total of 1794 incidents, representing 92.5% of the entire dataset. Table 1 contains the summary of incidents dataset. "# of Days" is the number of days where at least one incident was reported.

The IPS dataset does not include the case where a computer inside the organization attacks another computer inside the organiza-

Table 3: Regression results - how much does IPS metrics influence security incidents?

Variables	Model 1 (G1)	Model 2 (G1)	Model 3 (G2)	Model 4 (G2)
Al_t	-0.0002 (0.0001)	-	-0.0003 (0,0006)	-
At_t	0.0003 (0.0002)	-	-0.0038 (0.0154)	0.0292 *** (0.0090)
T_t	0.0025 * (0.0014)	0.0010 (0.001)	0.01517 ** (0.0059)	-
S_t	-0.0053 (0.0057)	-0.0061 (0.0057)	-0.0071 (0.0214)	0.0012 (0.0212)
AlT_t	0.2636 (0.1824)	-0.0284 * (0.0169)	-	-
AtT_t	-0.3031 (0.2326)	0.0154 (0.0330)	-	-
TAt_t	-	-	0.0047 (0.5398)	1.211 *** (0.27)
$AlAt_t$	-	-	-0.01645 (0.0348)	-0.0554 *** (0.0198)
	* p < 0.1	**$p < 0.05$	***$p < 0.01$	( ): standard error

Table 4: Regression results - Attacker's perspective

Variables	Bot (G1)	Bot (G2)	Hacking (G1)	Hacking (G2)	Spam (G1)	Spam (G2)
Al_t	-0.0001*** (0.00004)	0.00005 (0.0002)	-0.00007 (0.00008)	-0.0005* (0.00031)	-0.00005 (0.00012)	0.0002 (0.0005)
At_t	0.00015*** (0.00005)	-0.0013 (0.0051)	0.0001 (0.00009)	0.0181 ** (0.0084)	0.00004 (0.00014)	-0.01938 (0.0128)
T_t	0.0033*** (0.0004)	-0.00176 (0.0019)	0.00015 (0.0008)	0.009437 *** (0.003241)	-0.00089 (0.00113)	0.00851 * (0.0049)
S_t	-0.0004 (0.0016)	-0.0033 (0.0071)	-0.0015 (0.00335)	-0.00890 (0.01174)	-0.0033 (0.0046)	0.0059 (0.0178)
AlT_t	0.1768*** (0.051)	-	0.0872 (0.1067)	-	0.0450 (0.1467)	-
AtT_t	-0.1891** (0.065)	-	-0.1350 (0.1360)	-	-0.0346 (0.1871)	-
TAt_t	-	0.5390*** (0.1796)	-	-0.6828 ** (0.2966)	-	0.0701 (0.4503)
$AlAt_t$	-	-0.0287** (0.0115)	-	0.0316 (0.0191)	-	-0.0176 (0.0291)
	* p < 0.1	**$p < 0.05$	***$p < 0.01$	( ): standard error		

Table 5: Significant metrics per incident group

Variables	Bot	Hacking	Spam
Al_t (G1)	Sig. (***)	-	-
At_t (G1)	Sig. (***)	-	-
T_t (G1)	Sig. (***)	-	-
S_t (G1)	-	-	-
AlT_t (G1)	Sig. (***)	-	-
AtT_t (G1)	Sig. (**)	-	-
Al_t (G2)	-	Sig. (*)	-
At_t (G2)	-	Sig. (**)	-
T_t (G2)	-	Sig. (***)	Sig. (*)
S_t (G2)	-	-	-
TAt_t (G2)	Sig. (***)	Sig. (**)	-
$AlAt_t$ (G2)	Sig. (**)	-	-

tion. Table 2 contains the summary of IPS data. The "# of Attackers", "# of Targets" and "# of Signatures" represents respectively the cumulative number of attackers, targets and signatures.

The number of security incidents will be denoted as I_t, where t is the number of weeks and $t = 1, \ldots, 209$. The IPS metrics will be denoted as: number of alerts $= Al_t$, number of (distinct) attackers $= At_t$, number of (distinct) targets $= T_t$, number of (distinct) signatures $= S_t$, number of alerts per target $= AlT_t$, number of attackers per target $= AtT_t$, number of alerts per attacker $= AlAt_t$ and number of targets per attacker $= TAt_t$. In an effort to investigate how the IPS metrics might be linked to the number of security incidents, we built a multiple linear regression model. With a multiple linear regression model, it is possible to detect the effect of the independent variables on the dependent variable using a variable selection approach, that is, the screening of the candidate variables to obtain a regression model that contains the optimal subset of independent variables. Our dependent variable is the weekly number of incidents I_t and the independent variables are Al_t, At_t, T_t, S_t, AlT_t, AtT_t, $AlAt_t$ and TAt_t. The main idea is to select the independent variables, run the regression model and study its significance through the p-value obtained for each variable.

3.2 Results

We propose two regression models for each IPS metric group to investigate empirical evidence that IPS metrics might be linked to security incidents. In the first model, we analyze all G1 metrics, (Al_t, At_t, T_t, S_t, AlT_t and AtT_t). The second model is based on the analysis proposed by [3], which excludes the number of alerts and attackers. The authors state that the number of alerts and attackers might not be relevant, since an attacker could launch an attack against a target several times and at different times. This way, persistence is already handled by the number of alerts per target metric and number of alerts per attacker metric. We are therefore left with four metrics (T_t, S_t, AlT_t and AtT_t). The third model includes all G2 metrics (Al_t, At_t, T_t, S_t, $AlAt_t$ and TAt_t). Finally, the fourth model is also based on [3], which excludes the number of alerts and the number of targets (At_t, S_t, $AlAt_t$ and TAt_t). Table 3 summarizes the regression models and results. In all four regression models, at least one variable significantly impacts the number of security incidents. This result suggests that IPS metrics might be linked to the number of security incidents, and it can serve as a starting point to guide further empirical analysis between IPS metrics and the number of security incidents.

According to the regression results presented in Table 3, we found that four different metrics from G2 (number of targets, number of alerts, alerts per target and targets per alerts) significantly impact the number of incidents, while only two metrics from G1 (number of targets and alerts per target) significantly impact the number of incidents. The results show that the G2 metrics are more closely linked to security incidents. One possible reason might be the high number of false alerts raised by the metrics in the G1 group. However, the results could differ for various reasons, for instance, it might depend on the incident type. Table 4 summarizes the regression results for the three different incident groups: Bot, Hacking and Spam. We noticed that the significant metrics for each attacker perspective could differ according to the incident type. In other words, although we are analyzing the same incident dataset, the security metrics implementation process should take into account additional factors such as the attacker perspective.

Based on Table 4, Table 5 summarizes the significant metrics for three different incident groups: Bot, Hacking and Spam. The "Sig." represents a significant coefficient at (*) 0.1, (**) 0.05 and (***) 0.01. Table 5 shows that different types of security incidents have different significant metrics. Indeed, considering the G1 metrics, Bot was the only incident category with significant metrics. Similar behavior can be found for G2 metrics. The number of alerts (Al_t), for instance, was found to have a significant impact only on incidents Hacking.

4. DISCUSSION

The empirical analysis reveal three main results: i) IPS metrics

are linked to security incidents, ii) IPS metrics related to attacker's perspective are linked to security incident and iii) different types of security incidents have different significant metrics. These results can be used for identifying possible candidates for security incident indicators, developing methods to improve incident prevention and helping organizations interpret their IPS's better in the future. In the same way that other fields of science, such as economics, security indicators allow analysis of security performance and predictions of future performance. For instance, across all G1 results, we found that the number of alerts, number of attackers, number of targets, number of alerts per target and number of attackers per target significantly impact the number of incidents Bot. In other words, these metrics could be seen as likely candidates for security incident indicators and it can be used to perform analysis of security incident reporting process and to build more reliable security incident prediction models. Table 4 shows that the higher significant coefficients for incidents Bot are the number of alerts per target (0.1768), the number of attackers per target (-0.1891) and the number of targets per attacker (0.539). Thus, the number of alerts per target is associated with an increase of about 0.1768 incidents per week, the number of targets per attacker is associated with an increase of about 0.539 incidents per week and the number of attackers per target is associated with a decrease of about 0.1891 incidents per week. This result reveals that incidents Bot increases when there are many computers being targeted but from a small number of attackers (from outside) and when there are many external targets being attacked (from inside). A security analyst could use this information to create new prevention rules and to study trends and patterns related to security incidents.

Our study also indicates that IPS devices could be analyzed according to the degree of correlation found between the IPS metrics and the number of security incidents. Table 5, for example, shows that the investigated IPS device is more closely linked to incidents Bot. Additional research should be conducted for evaluating other IPS devices and understanding the correlation between IPS and different types of incident. This information might be useful for providing insight into what IPS's are particularly good for.

At last, our findings might be restricted to networks like those of universities: with nodes that are not fully controlled by the IT department. Private organizations, for instance, have different concerns about security. Therefore, a similar study, when conducted in such organizations, could show helpful but different results.

5. THREATS TO VALIDITY

Because of the method used by the Division of Information Technology to validate the incidents, we can assume that all incidents used in our work are real. Thus, there are no false positives among the incidents reported. However, we cannot quantify the number of undetected attacks and intrusions that did not lead to a security incident. The main issue with IPS event data is that the collected data are not perfect [3]. In other words, collected data might contain false positives and might not detect some malicious activity (false negatives). Besides, we cannot prove that a blocked attack would have been damaging to the targeted computer. In particular, for an attack to be successful, the targeted computer should have the associated vulnerability. As with all empirical studies, our results are limited to the datasets we investigated. In order to generalize our observations from this study to other environments, further studies should be performed.

6. CONCLUSIONS

In this paper, we investigated relationships between security metrics based on IPS data and computer security incidents, datasets collected at the University of Maryland. The results derived from our empirical model can be used for: identifying possible candidates for security incident indicators, developing methods to improve incident prevention and helping organizations interpret their IPS's better in the future. Future research should be conducted to compare the analysis for some other datasets and investigate the differences between them and also to evaluate the impact of security incidents over other variables, such as network topology and additional security incidents categories.

7. ACKNOWLEDGMENTS

The authors would like to thank the State of São Paulo Research Foundation (FAPESP) that supports this work. We also thank Gerry Sneeringer and the Division of Information Technology at the University of Maryland for allowing and supporting the described research.

8. REFERENCES

[1] M. Benini and S. Sicari. Risk assessment in practice: A real case study. *Computer Communications*, 31(15):3691–3699, 2008.

[2] W. Boyer and M. McQueen. Ideal based cyber security technical metrics for control systems. In *Critical Information Infrastructures Security*, volume 5141 of *Lecture Notes in Computer Science*, pages 246–260. Springer Berlin / Heidelberg, 2008.

[3] D. Chrun, M. Cukier, and G. Sneeringer. On the use of security metrics based on intrusion prevention system event data: An empirical analysis. In *HASE '08: Proceedings of the 2008 11th IEEE High Assurance Systems Engineering Symposium*, pages 49–58, 2008.

[4] M. Howard, J. Pincus, and J. Wing. Measuring relative attack surfaces. *Computer Security in the 21st Century*, pages 109–137, 2005.

[5] W. Jansen. Directions in security metrics research. Technical report, National Institute of Standards and Technology (NIST), 2009.

[6] A. Jaquith. *Security metrics: replacing fear, uncertainty, and doubt*. Addison-Wesley Professional, 2007.

[7] M. Swanson, N. Bartol, J. Sabato, J. Hash, and L. Graffo. Performance measurement guide for information security. Technical report, NIST Special Publication 800-55, 2003.

[8] The Center for Internet Security. The cis security metrics v1.1.0, November 2010.

[9] V. Verendel. Quantified security is a weak hypothesis: a critical survey of results and assumptions. In *NSPW '09: Proceedings of the 2009 workshop on New security paradigms workshop*, pages 37–50, 2009.

[10] C. Wang and W. Wulf. A framework for security measurement. In *Proc. National Information Systems Security Conference, Baltimore, MD*, pages 522–533. Citeseer, 1997.

[11] L. Wang, A. Singhal, and S. Jajodia. Measuring the overall security of network configurations using attack graphs. In *Proceedings of the 21st annual IFIP WG 11.3 working conference on Data and applications security*, pages 98–112. Springer-Verlag, 2007.

[12] J. Zalewski, S. Drager, W. Mckeever, and A. Kornecki. Can we measure security and how? In *Proceedings of the Seventh Annual Workshop on Cyber Security and Information Intelligence Research*, CSIIRW '11, 2011.