Cyber Situational Awareness through Network Anomaly Detection: State of the Art and New Approaches

Ivo Friedberg, Florian Skopik, and Roman Fiedler

Safety & Security Department
AIT Austrian Institute of Technology, Austria,
firstname.lastname@ait.ac.at

Abstract. With a major change in the attack landscape, away from well known attack vectors towards unique and highly tailored attacks, limitations of common rule- and signature-based security systems become more and more obvious. Novel security mechanisms can provide the means to extend existing solutions in order to provide a more sophisticated security approach. As critical infrastructures get increasingly accessible from public networks they show up on attackers' radars. As a consequence, establishing cyber situational awareness on a higher level through incident information sharing is vital for assessing the increased risk to national security in the cyber space. But legal obligations and economical considerations limit the motivation of companies to pursue information sharing initiatives. To support companies and governmental initiatives, novel security mechanisms should inherently address limiting factors. One novel approach, AECID, is presented that accounts for the limitations of many common intrusion and anomaly detection mechanisms; and which further provides the features to support privacy-aware information sharing for cyber situational awareness.

1 Introduction

Situational awareness [11] gets additional attention by governmental bodies as critical infrastructures and their security become part of national security plans. As these infrastructures get equipped with remote accessibility their cyber security aspects come into the spotlight of both: governmental bodies and potential adversaries. Various factors can currently be seen as limiting for an increased cyber situational awareness, the most prevalent in connection with information sharing. Cyber attacks carry the risk of reputation damage to affected companies what makes them reluctant to share relevant information timely. Various data formats to describe cyber incidents restrict the interoperability between companies. Finally legal restrictions specify restrictions on data that can be shared. This further introduces the need for sanitization of all data that is shared introducing additional delays. To increase cyber situational awareness novel security mechanisms are needed to enable companies to share information while complying to the law.

Since the emergence of the first ICT networks significant effort went into securing critical assets. But current security solutions also showed significant flaws when it comes to the prevention and detection of specific, novel and tailored attacks. Multi stage attacks leverage various social as well as technical techniques in multiple stages to breach a target's defenses. As an example Social Engineering attacks or Spear Phishing mails can be used to retrieve valid user credentials that might then be used for further system penetration on a more technical level using malware in a next stage. Such attacks are often unique to one target what makes them hard to detect for signature based security mechanisms. These mechanisms rely on knowledge about known attack patterns. By monitoring system behavior they can identify these patterns and notify operators or take actions to mitigate the ongoing attacks automatically. Signature based methods still play a crucial role in modern security systems due to their low false positive rates and effectiveness against known exploits. Anomaly based approaches try to fill the gaps of signature based mechanisms when it comes to novel attacks but introduce new limitations like high false positive rates.

Today's critical infrastructure providers can choose from a variety of technical security solutions that deal with different attack schemes on all levels: Firewalls that filter traffic at network borders, malware scanners that investigate binaries and executables for suspicious behaviour, intrusion detection systems (IDSs) that monitor events all over a network and verify them against predefined rules or anomaly detection methods to identify novel attacks. These systems are not only indispensable for the security of ICT infrastructures, but also provide additional information on the number, types and vectors of attacks a company is facing, and – when monitored over time – eventually help to establish cyber situational awareness.

In this paper, we therefore give an overview on cyber situational awareness and state-of-the-art anomaly detection approaches. With an example of a promising future system concept, called AECID - Automatic Event Correlation for Incident Detection [7] - we highlight how novel security mechanisms can be designed to address limiting factors of information sharing.

2 Cyber Situational Awareness

Information about threats, vulnerabilities and indicators of compromise is a valuable good for system administrators of today's complex and interconnected ICT systems. But sharing of this information is only one aspect of the greater concept called situational awareness. Different models of situational awareness where presented in [4],[6] or [15]. The most extensive of these models by Endsley ([4]) distinguishes three increasing awareness levels: *Perception* describes a level of awareness of the current situation. At the level of *comprehension* awareness about the reasons for the current situation is included. Finally, *projection* includes the ability to predict how the situation will evolve in the future. For the paper at hand we focus on the first level (*perception*).

The process of establishing situational awareness is currently mainly performed on an organizational level. Computer Emergency Response Teams (CERTs) are one example for an initiative across company borders. They collect and provide relevant security information and offer support for incident mitigation. But more needs to be done to effectively leverage the full potential for security incident mitigation. A number of standards where developed for information sharing to further establish an extensive situational awareness picture. Notable efforts resulted in standards by NIST [13], ITU-T [10] and ISO [9]. Central coordinating entities are proposed which collect reports on incidents, network monitoring data or status information of critical services by affected companies [8].

As mentioned before, limiting factors for information range from legal obligations to the risk of reputation loss. Examples for legal obligations come into play when talking for example about sharing of network monitoring data or log-files. Such data can contain sensitive user data (e.g. credentials, billing information) or classified company data which the company is not allowed to distribute outside the company's borders. One way of tackling this problem is to limit information sharing to the sharing of threat information. Threats in this case can be identified after incidents occurred and got analyzed locally. This information is then automatically sanitized. But this comes with a price: Threat information is often not received timely to the incident. Technical security solutions can and should therefore be designed with these limitations in mind to support information sharing.

3 Anomaly Detection

Anomaly detection is an actively researched field in many domains. Chandola et al. [3] identified the application domains as intrusion detection, fraud detection, medical and public health anomaly detection, image processing, anomaly detection in text data and anomaly detection in sensor networks apart from other not so prominent domains. In each domain we find different problems to solve, as well as different sets of data. For the paper at hand the most relevant application domain is

intrusion detection. An intrusion can be seen as an attempt to violate one of the three properties of the security triangle, namely availability, confidentiality and integrity. Intrusion Detection Systems (IDSs) aim at detecting those intrusions to take actions from triggering warnings to actively preventing the attacker from causing further harm. Literature as [18] or [14] classifies IDSs by different means. Anomaly detection can be one possible way for an IDS to identify intrusions.

Since anomalies are deviations of normal behavior, a system has to define a ground truth before anomalies can be detected. A learning phase is introduced prior to operation in which the anomaly detection system learns the expected behaviour of the data that should be analyzed. Two ways of learning normal behavior can be differentiated. In supervised learning, the training data consists of pairs. Each pair contains a sample and the expected result (anomalous or normal in this case). In unsupervised learning, the training data is unlabeled. It is the task of the anomaly detection tool to derive characteristics of the data from the unlabeled data set. After processing the training data, the goal of the anomaly detection approach is to make a decision (normal or anomalous) for new data samples. Training data is used to define what is to be considered normal behavior of a system [3,19]. This notion about normality is further used to mark patterns (also known as events or instances) in the data as normal or abnormal. It is important to note that normal behavior is not static but changes with the course of time. Therefore the notion of normality has to be periodically refined to represent the current system [19].

Chandola et al. [3] distinguishes three kinds of anomalies: **(i) Point Anomalies**, if a single event can be considered anomalous given the notion of normality we call it point anomaly. **(ii) Contextual Anomalies** when an event can be considered anomalous in respect to a given context. This contextual evaluation has to be encoded in the formulation of the problem. We can then deduce an anomaly given the events' behavioral attributes in its context. The same attributes might not be considered anomalous in another context. **(iii) Collective Anomalies**, if a series of events is considered anomalous we call it collective anomaly. Each event on its own in some other place in the stream might not be considered an anomaly. But the collective relation between them makes them anomalous.

Two main types of anomalies in ICT systems are described in [16]: (i) anomalies due to system failures and (ii) security related problems resulting in anomalies. Various classifications of anomaly detection approaches were taken by [3,19,18,16] just to name a few. The broadest classification by [3] distinguishes six classes with various subclasses from all of the before described domains: *classification based anomaly detection techniques*, *nearest neighbor-based anomaly detection techniques*, *clustering based anomaly detection techniques*, *statistical anomaly detection techniques*, *information theoretic anomaly detection techniques* and *spectral anomaly detection techniques*.

Zhang et al. [19] and Thottan et al. [16] distinguish anomaly detection approaches with focus on ICT networks. Zhang et al. [19] distinguishes *anomaly detection using statistics*, *anomaly detection using classifier*, *anomaly detection using machine learning* and *anomaly detection using finite state machines*. Thottan et al. [16] distinguishes between *rule-based approaches*, *finite state machines*, *pattern matching* and *statistical analysis*.

Various challenges in anomaly detection are identified by [3]. One of the severest challenge was already noted earlier: the definition of a complete notion of normal behavior. Fuzzy borders between normal and abnormal behavior and the fact that the notion of normality is evolving contribute to that fact. Further, for the case that the anomaly is caused by malicious behavior, it has to be considered that an attacker disguises his actions by making them look normal. In some cases it might be possible to tamper with the testing data; other times malicious actions are hidden in plain sight by the use of standard protocols (e.g. the Zeus malware [2] transmits captured data using encrypted payload in standard HTTP packets). Connected with the problem of getting a complete notion of normality is also the problem of getting representative, labeled training data and to differentiate noise from actual anomalies [1].

Large numbers of novel anomaly detection mechanisms in different variants and adapted to different use-cases where presented in recent years. [12,20,21,16,17] are just some examples to show the diversity in the approaches. This research trend is not likely to stop in the near future since anomaly detection is still seen as the way out of the limitations of signature based solutions when it comes to the detection of novel multi-stage attacks. Other legitimation for the ongoing research can be found in the limitations of today's anomaly detection mechanisms. Effective anomaly detection mechanisms are often very tailored to specific domains and application areas. Their set-up therefore includes a high customization and configuration effort that is often not feasible for large systems. Furthermore privacy concerns are often a limited factor as the generation of the normality notation often requires long-term storage of critical data. A good overview on anomaly detection in the domain of intrusion detection can be found in [18].

4 AECID - Automatic Event Correlation for Incident Detection

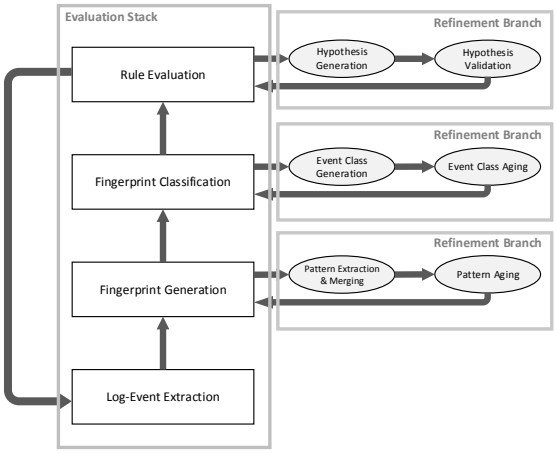

Fig. 1. Conceptual overview.

In order to address common issues of many of today's systems, we recently proposed a novel concept, AECID [7], that was designed to overcome previously mentioned shortcomings such as high manual configuration effort and legal issues regarding storage of network data containing sensitive data. AECID further supports the idea of cross-enterprise incident information sharing due to its data compression schema that effectively avoid privacy issues in such scenarios.

As part of the operational management of the ICT and industrial control systems in a critical information infrastructure, logging data is produced to report events, internal state changes, and committed actions. A conceptual overview of AECID is given in Fig. 1. First, AECID collects log files from systems in the ICT domain, maintaining the temporal order of log messages, and creates search patterns from single log lines. Logs can origin from systems like firewalls or routers but also from applications like a web server. A sample log line can be seen in Lst. 1.1. Second, AECID inspects all new log lines for those patterns and creates fingerprints for them depending on the presence or absence of individual patterns (see Tab. 1). This step significantly reduces the amount of data to handle and tremendously speeds up all further steps. It further obfuscates the sensitive information contained in the log lines. Since all further operations are performed on the abstract fingerprint it is not possible to draw conclusions about sensitive data without knowledge of the pattern set in the model.

```
1 service-3.v3ls1316.d03.arc.local apache: 2227 169.254.0.3:80 "mantis-3.v3ls1316.d03.arc.local" "
     mantis-3.v3ls1316.d03.arc.local" 169.254.0.2 - - [12/Feb/2014:13:30:16 +0000] "GET␣/mantis/
     login_page.php␣HTTP/1.1" 200 1343 "-" "Mozilla/5.0␣(X11;␣Linux␣i686)␣AppleWebKit/537.36␣(KHTML,␣
     like␣Gecko)␣Chrome/29.0.1547.65␣Safari/537.36"
```

Listing 1.1. Apache log excerpt from test environment.

Third, fingerprints and therefore underlying log lines, are classified based on the types of events that caused them. Forth, the aim of AECID is to correlate the different events, expressed in log lines, including their relative position to each other. For that purpose, it creates hypotheses about causes and effects. If a hypothesis can be confirmed as valid, it becomes a part of a model of the

Table 1. Example of a fingerprint. Consider the sample log-line in Lst. 1.1 from an Apache server running Mantis (https://www.mantisbt.org) in our test environment. This table shows an example set of patterns $\mathbb{P}'$ and how a fingerprint of the line in Lst. 1.1 would look like.

	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9
Patterns $\mathbb{P}'$:	GET	POST	[12/Feb/2014:13:30:15 +0000]	v3ls13	s1316.d0	ice-4.v3	apache:	login_page.php	mysql-n
Fingerprint:	1	0	1	1	1	0	1	1	0

system. These rules are continuously evaluated, extended and updated to build a dynamic notion of normality. An ongoing statistical analysis of correlation results allows the system to infer the degree of deviation of "if-then" hypotheses related to events, and therefore determine a degree of anomalous behavior.

Eventually, AECID provides additional benefits in multiple dimensions: (i) The model is generated following the real-world dependencies of components in the monitored environment. The detected relationships involve expected relationships (e.g. firewall rules being evaluated before a request is transmitted to a secured server) but are not limited to those. (ii) The rules in the model are automatically generated and unique for each monitored environment. A potential attacker cannot easily tailor an attack to prevent rules from failing; the attack, after all, alters the system behavior. (iii) Syntax and semantics of the log lines are widely irrelevant, since the model is tailored to the log input and the pattern creation and selection purely based on statistical analysis. This prohibits attacks on the analyzed data since an attacker cannot determine wich information is leveraged by the system. (iv) Sharing of sensitive information about security incidents with state actors or other third parties is usually forbidden by companies, although information sharing is a requirement to effectively monitor critical infrastructures on a national level, e.g., as demanded by the EU NIS directive [5]. Given the abstract form of information AECID transforms log lines to, information sharing is enabled without compromising privacy.

AECID analysis host based data from various components in a network in order to make assumptions about relations between the components on the network level. Therefore, the approach can be considered as an *Information Theoretic Anomaly Detection Technique* based on the clasification presented in Sect. 3. From the way the data sources are analyzed and evaluated it is also a *Statistical Anomaly Detection Technique* with a parametric characteristic. Learning works unsupervised.

5 Conclusion

This paper provided a short overview about the state-of-the-art in the field of situational awareness and anomaly detection. Different characteristics and classification methods for anomaly detection systems where presented. Although considerable research is undertaken with respect to ICT security mechanisms and situational awareness, this article showed that there are still limiting factors that prohibited the development of generically effective solutions.

The paper then outlined a proposal of a novel anomaly detection approach, called AECID, that aims at extending existing intrusion detection systems. AECID tackles many of the limitations of previous research results. Since it is not dependent on specific systems or domains, it provides a generically effective security solution. By abstracting from sensible information in the analyzed data at an early processing stage, it is especially suited for privacy-aware incident information sharing. This form of cooperation between critical infrastructure providers gets increased attention by governmental bodies in the course of national cyber security strategies. Here, AECID is specifically designed to support the establishment of cyber situational awareness.

Acknowledgements

This work was partly funded by the Austrian FFG research program KIRAS in course of the project CIIS (840842) and the European Union FP7 project ECOSSIAN (607577).

References

1. Václav Bartoš and Martin Žádník. Network anomaly detection: comparison and real-time issues. In *Dependable Networks and Services*, pages 118–121. Springer, 2012.
2. Hamad Binsalleeh, Thomas Ormerod, Amine Boukhtouta, Prosenjit Sinha, Amr Youssef, Mourad Debbabi, and Lingyu Wang. On the analysis of the zeus botnet crimeware toolkit. In *Privacy Security and Trust (PST), 2010 Eighth Annual International Conference on*, pages 31–38. IEEE, 2010.
3. Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Computing Surveys (CSUR)*, 41(3):15, 2009.
4. M.R. Endsley. Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1):32–64, 1995.
5. European Commission. Commission proposal for a directive concerning measures to ensure a high common level of network and information security across the union, 2013. http://ec.europa.eu/digital-agenda/en/news/commission-proposal-directive-concerning-measures-ensure-high-common-level-network-and.
6. M.L. Fracker. Measures of situation awareness: Review and future directions. Technical Report AL-TR-1991-0128, Wright-Patterson Air Force Base, 1991.
7. Ivo Friedberg, Florian Skopik, Giuseppe Settanni, and Roman Fiedler. Combating advanced persistent threats: From network event correlation to incident detection. *Computers & Security*, 48:35–57, 2015.
8. Jorge L Hernandez-Ardieta, Juan E Tapiador, and Guillermo Suarez-Tangil. Information sharing models for cooperative cyber defence. In *Cyber Conflict*, pages 1–28, 2013.
9. ISO. Iso/iec27010: Info. tech.: Security techniques - information security management for inter-sector and inter-organizational communications. 2012-03-20.
10. ITU-T. Recommendation itu-t x.1500 cybersecurity info. exchange tech., 2012.
11. Sushil Jajodia, Peng Liu, Vipin Swarup, and Cliff Wang. *Cyber Situational Awareness: Issues and Research*. Springer, 2009.
12. Guichong Li, Nathalie Japkowicz, and Lian Yang. Anomaly detection via coupled gaussian kernels. In *Advances in Artificial Intelligence*, pages 343–349. Springer, 2012.
13. NIST. Framework for improving critical infrastructure cybersecurity. 2014-02-12.
14. F Sabahi and A Movaghar. Intrusion detection: A survey. In *Systems and Networks Communications, 2008. ICSNC'08. 3rd International Conference on*, pages 23–26. IEEE, 2008.
15. N.D. Sarter and D.D. Woods. Situation awareness: A critical but ill-defined phenomenon. *International Journal of Aviation Psychology*, 1:45–57, 1991.
16. Marina Thottan and Chuanyi Ji. Anomaly detection in ip networks. *Signal Processing, IEEE Transactions on*, 51(8):2191–2204, 2003.
17. Jian Yin, Gang Zhang, Yi-Qun Chen, and Xian-Li Fan. Multi-events analysis for anomaly intrusion detection. In *Machine Learning and Cybernetics, 2004. Proceedings of 2004 International Conference on*, volume 2, pages 1298–1303. IEEE, 2004.
18. Yingbing Yu. A survey of anomaly intrusion detection techniques. *Journal of Computing Sciences in Colleges*, 28(1):9–17, 2012.
19. Weiyu Zhang, Qingbo Yang, and Yushui Geng. A survey of anomaly detection methods in networks. In *Computer Network and Multimedia Technology, 2009. CNMT 2009. International Symposium on*, pages 1–3. IEEE, 2009.
20. Ya-ling Zhang, Zhao-guo Han, and Jiao-xia Ren. A network anomaly detection method based on relative entropy theory. In *Electronic Commerce and Security, 2009. ISECS'09. Second International Symposium on*, volume 1, pages 231–235. IEEE, 2009.
21. Ying Zhao, Zhigao Zheng, and Hong Wen. Bayesian statistical inference in machine learning anomaly detection. In *Communications and Intelligence Information Security (ICCIIS), 2010 International Conference on*, pages 113–116. IEEE, 2010.