The Seven Golden Principles of Effective Anomaly-Based Intrusion Detection

Florian Skopik, Markus Wurzenberger, and Max Landauer | Austrian Institute of Technology

The MITRE ATT&CK framework counts 530 ways to exploit enterprise systems—and every month new techniques are added. Cybersecurity vendors continuously offer new detective solutions, but purchasing, deploying, and maintaining a specific product is expensive. It's time to reflect on the underlying principles of effective anomaly-based intrusion detection.

"Prevention is ideal, but detection is a must." Active monitoring and intrusion detection systems (IDS) are the backbone of every effective cybersecurity framework. Whenever carefully planned, implemented, and executed preventive security measures fail, IDS are a key component of the last line of defense. IDS are an essential means to detect the first steps of an attempted intrusion in a timely manner. This is a prerequisite to avoid further harm. Security experts agree that active monitoring of networks and systems and the application of IDS are a vital part of the state of the art. Usually, findings of IDS as well as major events from monitoring are forwarded to, managed, and analyzed with security information and event management (SIEM) solutions.[1] These SIEM solutions provide a detailed view on the status of an infrastructure under observation.

However, a SIEM solution is only as good as the underlying monitoring and analytics pipeline. IDS are an inevitable part of this pipeline, which spans from gathering systems' data, including operating system logs, process call trees, memory dumps, and so on, feed them into analysis engines, and report findings to SIEMs. Obviously, the verbosity and expressiveness of data are a key criterion for the selection of data sources and associated analysis approaches. This is an art of its own and mainly dependent on answering what kind of common attack vectors today (referring to the MITRE ATT&CK matrix[2]) are reflected best in which sources [e.g., Domain Name System (DNS) logs, netflows, syscalls, and so on]. There are literally hundreds of tools and agents to harness the different sources and tons of guidelines on the configuration of these tools to control the verbosity and quality of resulting log data.

In terms of detection mechanisms, most commonly used today are still signature-based network intrusion detection system (NIDS) approaches. Similarly, signature-based host-based intrusion detection systems (HIDs) are capable of using host-based sources, such as audit trails from operating systems, to perform intrusion detection. The secret of their successes lies in their simple applicability and the virtually zero false positive rate. Either a malicious pattern is present, or it is not. It's as simple as that.

Unfortunately, this easy applicability comes with a price. The slightest modification to the malware or attacking tools changes the traces that an attack leaves on a system, which hinders the effectiveness of signature-based approaches. For instance, Christiansen[3] demonstrated that well-known malware can evade IDS by implementing a single NOP instruction in the right place of its code.

Digital Object Identifier 10.1109/MSEC.2021.3090444
Date of current version: 9 July 2021

To mitigate attacks with polymorphic and customized tools, IDS vendors combine signature-based approaches with heuristics to enable a kind of fuzzy detection, i.e., detect patterns that match to a certain degree but allow some inherent noise. This again, however, increases the false positive rate, which limits the use of such approaches. The job for security solution vendors and integrators is to find the sweet spot where fuzzy signature-based matching still works without producing too many false detections. However, even then, attackers that "live off the land"—in other words, who use just system tools that they find on the target systems—cannot be detected at all. For instance, if someone steals legitimate credentials, logs into a web-based platform, and starts to copy off data, no malware is ever used. This is one of the reasons why the focus on the detection of known bad actions seems to be a dead end for defenders in the long run.

As a consequence, a major transition away from signature-based approaches to behavior-based approaches takes place.[4] The fundamental idea is that if it is not possible to characterize what malicious activities look like to search for those in an infrastructure, the new aim should be to model or learn legit activities and treat everything else as potentially hostile. This is how anomaly detection (AD) methods work.

Threat hunting[5] complements this focus on advanced highly automated detection and puts the human operator in the center whose goal is to proactively search through infrastructures to discover and isolate adversarial actions that evade carefully planned security mechanisms. To achieve that, threat hunting applies anomaly-based intrusion detection in the form of user behavior analytics (UBA), among other tools. Regardless of whether anomaly-based IDS are applied as highly automated detective means to secure an infrastructure or used to proactively sift through data streams to semiautomatically isolate advanced adversarial actions, it is a vital key technology to secure our networks.

However, the pity of AD is that it is quite error-prone, verbose, and unspecific. This means that in a medium-sized enterprise system, depending on how tightly it is monitored and how strictly the normal behavior is defined, possibly hundreds of anomalies arise every minute. Obviously, this is by no means a feasible approach.

We need additional measures to compensate for this drawback and to make anomaly-based intrusion detection more appealing. In this article, we investigate numerous aspects of intrusion detection and threat hunting that can mitigate this problem and help to make modern anomaly-based intrusion detection applicable.

Key Questions of Anomaly-Based Intrusion Detection and Threat Hunting

Defending a complex infrastructure against adversaries is hard. Often, systems are not designed top-down but rather organically emerge driven by ever-growing business dynamics. Consequently, no one fully understands how the whole system really works; only specialists in certain areas are familiar with some isolated parts. Furthermore, even primitive components offer a multitude of configuration options and operating modes, which dramatically increases the attack surface and thus the way to exploit them in unprecedented ways. Effective anomaly-based intrusion detection is not a job for a "drop and go" solution. It needs attention to detail and continuous care taking—and eventually a structured approach to get it done right.

Numerous challenges of AD have been studied in recent years.[6] Most of them are centered on the fact that it is hard to define a baseline to compare against, that this baseline is a moving target, deviations are hard to interpret, or that only limited data are available to train machine learning methods. While these concerns are all true in principle, we argue that specifically for the cybersecurity domain, there are ways to relax the situation. Multiple options exist to support AD methods for cybersecurity that, if followed carefully, make these methods more applicable—both for automated detection of adversarial actions or even more complex threat hunting.

We started our investigation with a literature survey with the keywords "intrusion detection," "active defense," "cyber threat intelligence," and "threat hunting" in IEEE *Xplore*, the Association for Computing Machinery (ACM) Digital Library, Scopus, and Google Scholar. The aim was to identify the requirements on and common aspects of modern effective intrusion detection and threat-hunting approaches. Since we discuss the application aspects of modern intrusion detection in operational environments in this article, we did not conduct our survey with common academic sources alone but instead further focused on industry reports, white papers, and the collected wisdom of industry experts. Numerous well-known annual reports are available in the domain, including those from Mandiant/FireEye,[7] Crowdstrike,[8] the Ponemon Institute,[9] and SANS,[5] just to name the well-known ones. We studied those in detail and extracted common topics they address and common questions they ask to identify the important aspects of modern cybersecurity detection and response solutions and services. Additionally, we considered ongoing discussions from the cyberthreat intelligence-sharing domain, mainly OASIS Cyber Threat Intelligence (CTI). Following this approach, summarized in Figure 1, we ensured that we would not leave out any relevant point of view on anomaly-based intrusion detection.

Based on the survey results, we structure our investigations along seven key questions (see Figure 2) and by answering them come up with seven key principles that discuss important mechanisms to enable anomaly-based intrusion detection and threat hunting.

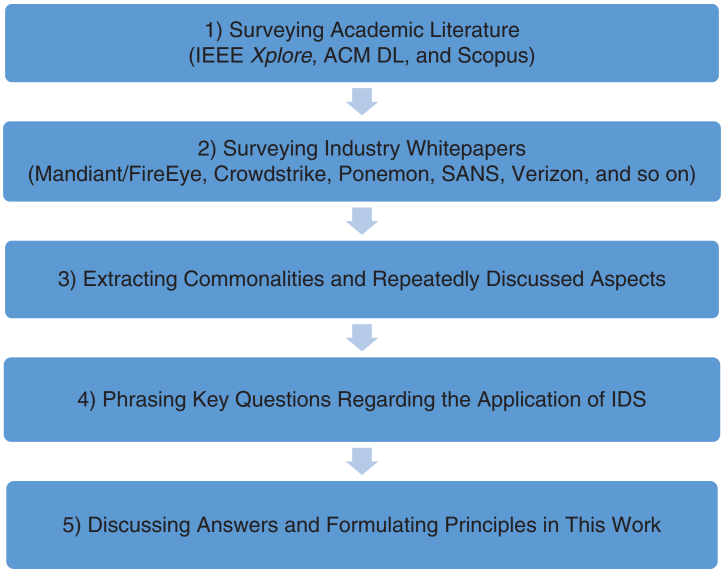
1) Surveying Academic Literature
(IEEE *Xplore*, ACM DL, and Scopus)
2) Surveying Industry Whitepapers
(Mandiant/FireEye, Crowdstrike, Ponemon, SANS, Verizon, and so on)
3) Extracting Commonalities and Repeatedly Discussed Aspects
4) Phrasing Key Questions Regarding the Application of IDS
5) Discussing Answers and Formulating Principles in This Work

Figure 1. A methodology to discover and formulate key principles.

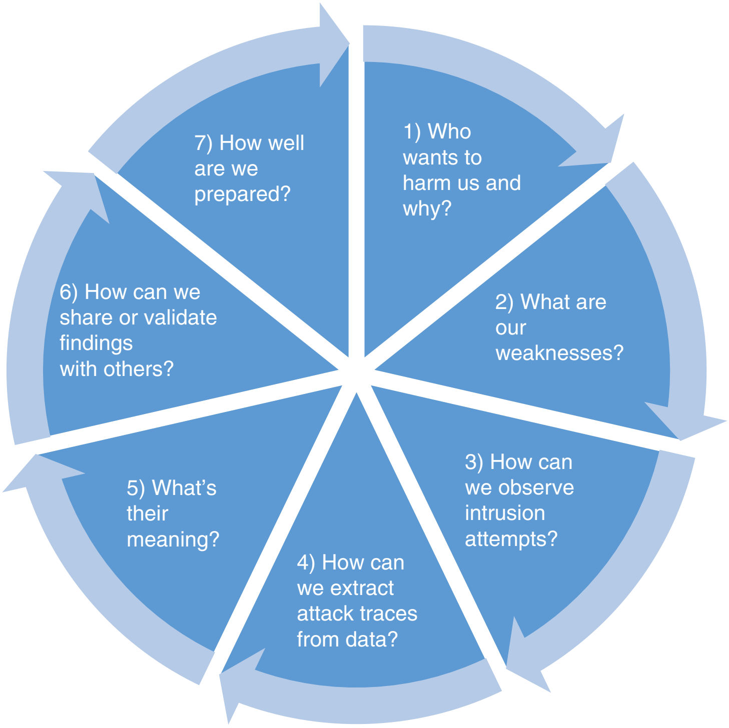
1) Who wants to harm us and why?
2) What are our weaknesses?
3) How can we observe intrusion attempts?
4) How can we extract attack traces from data?
5) What's their meaning?
6) How can we share or validate findings with others?
7) How well are we prepared?

Figure 2. Key questions of effective intrusion detection and threat hunting.

Each of these principles makes anomaly-based intrusion detection easier to apply and less error-prone. Figure 2 highlights these key questions raised when building an effective solution and underpins that this endeavor does not just have a technical dimension but requires a sound multidimensional methodological approach. Starting with the question of who wants to harm us and the motivation of the threat actors, it is important to match their motivation and capabilities to our own vulnerabilities. Only when these three factors—actor motivation and capability as well as our own vulnerabilities—are combined does a real threat for a successful intrusion exist. Once we expect intrusions to happen, the next question is to investigate how these might be observed, such as what log files or data streams can point to the presence of an adversary in our network. Of course, it is important to isolate attack traces in massive amounts of data with certain reliability and interpret their meaning correctly. After achieving this, we should further think about how we can share our findings with others in the same sector to warn them about ongoing campaigns in a timely manner—and in return get validation of our findings. Eventually, discovering any loopholes in our defense strategy and estimations on how well we are prepared for the next wave is an essential step of continuous improvement.

Guidance to Effective Anomaly-Based Intrusion Detection and Threat Hunting

We address these questions with the following guiding principles and shed some light on related aspects.

Principle 1: Get to Know Your Enemy

What is true for the real world also applies to the cyber domain. Knowing the motivation and capabilities of adversaries is key to an effective defense. This knowledge is equally important to set up effective monitoring as well as to timely spot the different forms of malicious activities. However, for many years, defenders were always one step behind as they needed to detect new attacks first, analyze their effects on the system (such as dropping new malware or changing system files), and then build signatures for malware detection and antivirus solutions. So-called indicators of compromise (IoCs) were important to verify that a system has been penetrated; however, simple IoCs, including file hashes, process names, or certain memory patterns, are easy to circumvent by slightly adapting attacker tools or obfuscating attack techniques. The new hot topic is therefore modeling of more complex tactics, techniques, and procedures (TTPs), which represent the adversaries' modus operandi and are harder to change, and thus are a more sustainable means for detection.

A further challenging research area is attribution, which aims to develop solutions for associating observed cyberattacks to well-known threat actors just by investigating their actions along the cyberkill chain[10] and deriving their likely objectives. Understanding the motivation and capabilities of the numerous advanced persistent threat groups, nation-states, and hacktivist is important. They all carry out targeted highly complex operations, albeit mainly toward high-profile targets such as governments and large industries. Eventually, we must not underestimate the threat stemming from botnets and opportunistic attackers that simply exploit a vulnerability because it is there. The cyberarsenal of this largest group of attackers is studied by researchers and nation-states using honeypots and honeynets[11] and poses vital knowledge to plan our defense strategies: specifically, to determine where to look for compromise.

Principle 2: Get to Know Thyself

Knowing the enemies and their nasty tricks is just half of the story. Equally important is to know the environment, its vulnerabilities, and its weaknesses that constitute its attack surface. The latter is highly influenced by the maturity level of the organization, determined by policies, standards, guidelines, and procedures that shape the behavior of all entities—being machines or humans. This is a mandatory prerequisite for deriving a baseline of good (and predictable) behavior. Unfortunately, it is also one of the most often neglected aspects of intrusion detection. The truth is, an organization whose maturity level is not advanced enough can never effectively apply anomaly-based intrusion detection or threat hunting and therefore is unable to capitalize on the advantages of modern machine learning-based intrusion-detection approaches.

If a policy or standard dictates what is allowed in an environment (and we assume that people will mostly stick to it), it is easy to spot deviations from this defined good behavior. Let's take a look at a couple of examples. A company standard may stipulate that Windows hosts are all present in a certain network segment, that they are just allowed to use a preconfigured web proxy to go to the Internet, that only the internal DNS may be used to query domains, or that only USB sticks of a certain vendor may be used. Of course, many of these requirements can be technically enforced, while others need more cooperation from employees. Adversaries may try to circumvent some of these enforced design decisions, for instance when they laterally move from one Windows host to another, although this is not foreseen by design. Such behavior is easily detectable as an anomaly if we know that it is not part of our normal and accepted baseline.

Besides the desired behavior defined in policies, guidelines, and procedures, there are other observable behavioral patterns that no one dictates top-down but that rather emerge bottom-up. For instance, some employees might show some steady work patterns, as they usually arrive at the same time in office and follow a rather deterministic work schedule. Other patterns are more technical in nature. An example for that is network link graphs that can be constructed by observing pairs of communicating machines, how much data they usually exchange in which direction, and what specific protocols they use for this. It is quite a surprise to find out how steady the behavior of a smoothly running business can be.

Gaining insights into our normal business conditions is part of modern intrusion detection. Recent attacks have impressively demonstrated that skilled adversaries stay below the radar by "living off the land," which means they utilize only unsuspicious operating system tools, which they find on target systems, such as powershell, wmic, or Server Message Block (SMB). They do not introduce any of the well-known adversarial tools. Thus, it becomes very hard—if not impossible—for traditional IDS to spot them. However, adversaries eventually utilize a system differently than legitimate users and do not stick to normal work conditions, for instance when they try to copy large volumes of data or access many files in a short time frame.

Obviously, talking about anomaly-based intrusion detection and threat hunting is pointless if there is no clear understanding of how a system is designed and usually utilized, and if its acceptable behavior is not defined. An applicable and enforced set of policies, standards, and procedures, as well as a clear system documentation in form of managed asset databases and configuration management databases (CMDB), a professionally maintained infrastructure, and a clear organizational structure reflected by distinct roles and their responsibilities, is good security practice, aids to a proper baseline, and is thus the underlying foundation of effective intrusion detection.

Principle 3: Be Open to All Sorts of Data

Researchers and security analysts have argued for some time about what type of data are best suited for intrusion detection. There are discussions centered on whether data captured from networks or host-based data offer more visibility; others argue over whether generic low-level kernel- and system-log data are more valuable than application-specific transaction records. Some focus on behavior analytics of processes, others on memory inspection and pattern mining. The truth is that in principle there is no right or wrong type of data. The feasibility of data to spot intrusions solely relies on the attack technique applied. The MITRE ATT&CK framework does a great job summarizing the wide variety of attack techniques and associated data sources to spot them. At the time of writing (April 2021), MITRE

ATT&CK lists 530 different forms just for standard enterprise environments. Taking a closer look and focusing on those techniques that are applicable in our own environment and individually rating their effects on our systems is a big step forward toward a deeper understanding of individual cyberrisks.

Today's SIEM use cases are usually centered on well-known attack cases. We need specific monitoring data to spot cross-site-scripting attempts on a web server, which is entirely different from discovering brute force login attempts against a remote shell. SIEM use cases help us to keep focusing on the essential attacks and carefully plan what data we shall ingest first. While this approach quickly leads to first successes, it creates problematic blind spots in the long run. Attacks that we do not anticipate are not covered by these use cases. Anomaly-based intrusion detection is a means to mitigate these issues, which may also be applied additionally to established signature-based approaches. Since there are so many techniques on the one side, and motivated attackers will quickly adapt and select those techniques that are most successful to penetrate our environment on the other side, it is best to not rule out any kind of data as being useful. Eventually, we need to be open to ingest all possible sorts of data, not just what our security consultants tell us is important today or what vendors think might be sufficient to implement. Be aware that not only raw data are of interest; further metadata such as an event's context (e.g., the time of day it was raised), its frequency within a defined time span, and surrounding elements are equally useful to characterize a system's behavior.

While it is important to not rule out the value of any kind of data, it is also true that data hoarding without use is poor advice. We must distinguish between tactical data that is carefully selected and sparse in nature and compliance data, the purpose of which is documentation through full coverage of past activities. Both are important but used in entirely different ways. While tactical data helps us to quickly and efficiently spot anomalous activities, compliance data are indispensable to further analyze a given situation and deduct potential consequences. The latter is a capability that is specifically important for incident response. For that reason, we should collect both tactical and compliance data but in different systems, used by people in different roles for different purposes.

Principle 4: Analyze Smart

Following the first three principles of this article, we should have collected all of the data that is necessary to spot an intrusion. However, intrusions do not serve up themselves on a silver plate. In a first step, we need to extract the relevant data points and at the same time reduce the potentially massive amounts of data in a smart way to be able to handle them properly. Feature extraction[12] is the science to do exactly that.

A wide variety of analysis techniques exist today and, depending on which ones we apply, we gain different sorts of anomalies and insights into a system's behavior. For instance, long-tail analysis focuses on identifying rarely occurring events, as does outlier detection based on clustering approaches. Other analysis techniques are time-series approaches, such as the autoregressive integrated moving average (ARIMA) model, used to spot long-term deviations in trends, as well as frequency detection, detection of event correlations, or changes in the value distribution of data fields. The art is to apply the right analysis technique on a matching set of data properties.

While we can always manually define what data we like to analyze with which kind of technique, numerous machine learning and artificial intelligence (AI) concepts have been proposed[6] in recent years to avoid this manual step: artificial neural networks, decision trees, support-vector machines, and Bayesian networks, just to name a few. In general, three classes of machine learning approaches can be distinguished: supervised, semisupervised, and unsupervised methods, depending on whether and to what extent they need to be trained with labeled data. Unsupervised learning does not require any labeled data and is able to learn to distinguish normal from malicious system behavior during the training phase. Based on the findings, it classifies any other given data during the detection phase. Semisupervised learning implies that the training set only contains anomaly-free data and is therefore also called "one-class" classification. Supervised learning requires a fully labeled training set containing both normal and malicious data.

All of these techniques, preferably used in combination, and many more, may spot deviations from a defined or learned baseline and may alert security analysts. However, in practice another problem arises: alert fatigue. Too many alerts in too-short time intervals causes security analysts to ignore alerts and therefore not handle them with the required attention or simply become overloaded.

Attempts to mitigate this problem include the aggregation of anomalies, for instance occurring within a certain time frame, their interpretation with the help of further context and prioritization to report only highly ranked detection results. For example, it is of far more interest to us that a system that deals with highly classified data or that supports one of our main business processes behaves odd than occurrences of odd events in any other part of the infrastructure. Anomaly rating and ranking requires a deeper understanding of the importance of potentially affected systems for our business. Incorporating asset inventories, CMDBs, and all further sources of contextual data help to determine the risks associated with the occurrence of anomalies.

Principle 5: Making Sense Out of Data

With signature-based solutions, which look for known bad behavior, it is simple to alert on observed specific malicious behavior. Usually, we then immediately know what we are up to, e.g., whether there is a backdoor in our system, a virus that modifies files, or a ransomware that encrypts contents. Unfortunately, it is not that easy with anomaly-based intrusion detection. Having spotted an anomaly or, rather, a set of anomalies, the next step is to derive its possible root cause. Numerous methods exist to achieve that—automatic ones applicable to IDS or semiautomatic ones often applied by security analysts in course of threat-hunting activities.

First, we can derive some useful information simply from knowing the (type of) affected system, the type of anomaly reported, or the specific data source that emitted the anomalous data. Matching this data with contextual information, such as criticality and sensitivity levels given in our asset database, already provides first insights.

Second, another way to gain insights is to make a lookup in historic data if the reported kind of anomaly (and pattern of multiple anomalies across different systems and data sources respectively) was observed in the past. Furthermore, reports from previous incident response activities might be of additional support. For instance, if we investigated unusual connection attempts between Windows clients in the past and found out that there was a malware that tried to compromise hosts in the vicinity via SMB, this is likely the same reason if we suddenly observe TCP traffic on port 445 with similar properties in the internal network. Matching observed behavior, which is not in accordance with our baseline, to historic situations usually works quite well for attack techniques that we have observed and investigated in more detail before.

However, an attack or malware exploiting one of our vulnerabilities the first time will not be revealed using this method. Here, which is the third case, we would request help from our peers to find out if they are familiar with our findings, before we would kick off our internal incident response process and perform time- and resource-intensive deeper investigations. The latter is the case if neither historic data nor information from peers can help us to understand an anomaly because it is either quite specific to our infrastructure or uses a zero-day vulnerability that no one knew about so far.

Before detected anomalies trigger a deeper investigation, a triage is advisable. This means that we carefully decide whether a finding justifies further investigations and, in case of multiple concurrent findings, decide on which ones we should primarily focus. Every incident investigation is an opportunity to learn and every activity a means to extend our knowledge. Nevertheless, we need to manage our scarce resources carefully and focus on those events that are truly of importance. Not every scanning attempt on an externally facing interface is worth detailed investigation.

Principle 6: Sharing Means Caring

Information sharing has gained tremendous momentum in recent years. Several initiatives, spanning from the automatic distribution of IoCs and sharing of suspicious artifacts to sharing knowledge in form of reports, have emerged. The motivation for sharing is usually twofold. First, by contributing knowledge of new attack vectors, ongoing malicious activities, and investigation results to the community, we hope to help others who may return the favor in the future. Second, as a side effect of sharing, we also seek validation of our findings. If we drop suspicious artifacts into a community and a handful of members confirm having discovered the very same, we immediately know that we have truly discovered something that justifies further investigations (i.e., we have minimized the probability of a false positive). Furthermore, we also know that we are not victim of a targeted attack. Overall, sharing allows us to collaborate with others in a likely similar situation and share investigation results and split resources spent.

Unfortunately, it is not as simple as stated here. A major hurdle is that everyone's systems and infrastructures are different, there is a multitude of potentially suspicious activities and artifacts, and if the result of IDS is truly a real positive ultimately depends again on our individual (!) baseline. The same behavior might be problematic for one organization but perfectly normal for another one. Furthermore, privacy and data sovereignty issues hinder information sharing as well.

From a research perspective, two major hurdles need to be taken. First, the organization that likes to share information must do so in a normalized form, understandable and usable for the receiver. The receiver, on the other side, needs to contextualize this normalized information, which means to interpret the received information for the specific organizational context. This task is not straightforward and needs to account for an organization's baseline, risk profile, known vulnerabilities, and some more aspects.

This is also the reason why automated information sharing with respect to both generation and consumption at a large scale does not work out well yet. A piece of information has likely different meaning for the sender and the receivers, as they do not know each other's organizational context and the circumstances in which the shared information was created or is being consumed.

In the simplest form, a semiautomatic sharing mechanism uses IoCs such as the malware information sharing platform (MISP).[13] Sharing IoCs is easy this way, yet ineffective; sharing TTPs, however, is hard and not well understood yet. With increasing standardization in this area, such as OASIS Structured Threat Information Expression (STIX),[14] we expect to see more advanced solutions in the near future.

The fine art to achieve is the sharing of knowledge about concrete attack tactics used by adversaries, combined with the immediate applicability of this knowledge in diverse environments, so-called actionable cyberthreat intelligence (CTI). Knowledge shared this way is vital to either tune detection engines or to proactively harden our own systems and improve baselines. Table 1 depicts an overview of the information being shared today at the different abstraction levels. All of these entities are part of the STIX model.

Principle 7: Learn and Prepare for the Next Wave

Learning from previous incidents, whether one's own or those of others, and deriving effective countermeasures to avoid similar problems in the future is a key principle for increasing the own security posture. With respect to anomaly-based intrusion detection and threat hunting, we should always raise the question of how we could have done better. Whether some monitoring needs to be adapted, new data sources integrated, the machine learning algorithms tuned, or our capabilities of interpreting signs of intrusions increased, an objective assessment of incidents is the basis for that.

However, we should not wait until something happens and assume that we are safe otherwise. Chances are high that our system is already compromised but we haven't recognized it yet. That is the reason why we should regularly benchmark the current configuration of the IDS and challenge it with new attack vectors. This could happen in a production system but might not always be the best option. Test setups, testbeds, and simulation environments are other means to tune configurations of IDS or try out new detection algorithms. Eventually, in times of freedom, we must ensure that our IDS approach is ready and effective.

Besides the ability to detect new variants of attacks, it is similarly important to carefully tune our exceptions. If we confuse legitimate virtual private network (VPN) connections with long-lived connections of botnets, if bandwidth peaks during the weekly backup put us on alert, our intrusion detection approach needs further tuning to be taken seriously. False positives must lead to carefully designed exceptions in our anomaly-based intrusion detection approach to get rid of any noise that may degrade the trustworthiness in the solution.

The most straightforward way to test the readiness of our intrusion detection approach, including both the whole technology stack and the people maintaining it, is to use red teaming.[15] Here, a professional company applies TTPs as real threat-actor groups do and try to penetrate a target system. In contrast to well-known penetration tests, the primary goal is not to discover and report technical vulnerabilities but to test the efficiency and effectiveness of IDS, their configurations, the way people handle discovered traces of intrusions, and the feasibility of the whole incident response process.

Table 1. Information being shared today on the different abstraction levels.

Levels	Artifacts	Questions answered
Strategic	Threat actors	Who is carrying out the operation?
	Campaigns	What is their motivation and goal?
	Tactics, techniques, and procedures (TTPs)	What is their modus operandi?
Operational	Incidents	Who was affected and how?
	Course of action (CoA)	What are countermeasures?
	Exploit target	What vulnerabilities do they exploit?
Tactical	Indicators	What general IoCs should we look for in our networks?
Atomic	Observables	What have we found, where and when (sightings)?

Effective anomaly-based intrusion detection needs to account for many aspects to make it applicable in real enterprise environments. State-of-the-art signature-based solutions are an important means of every basic defense, but elevated security levels are achievable only in a mature and professional environment, applying prudent monitoring, machine learning supported detection mechanism, and proper handling of results. This endeavor does not only call for a feasible technical solution but requires a sound multidimensional approach. We conducted a literature survey and summarized all of the aspects discussed in the major industry reports. Figure 3 associates each of the discussed principles with a certain part of an overall intrusion detection methodology for an enterprise environment. No single one of these principles is more

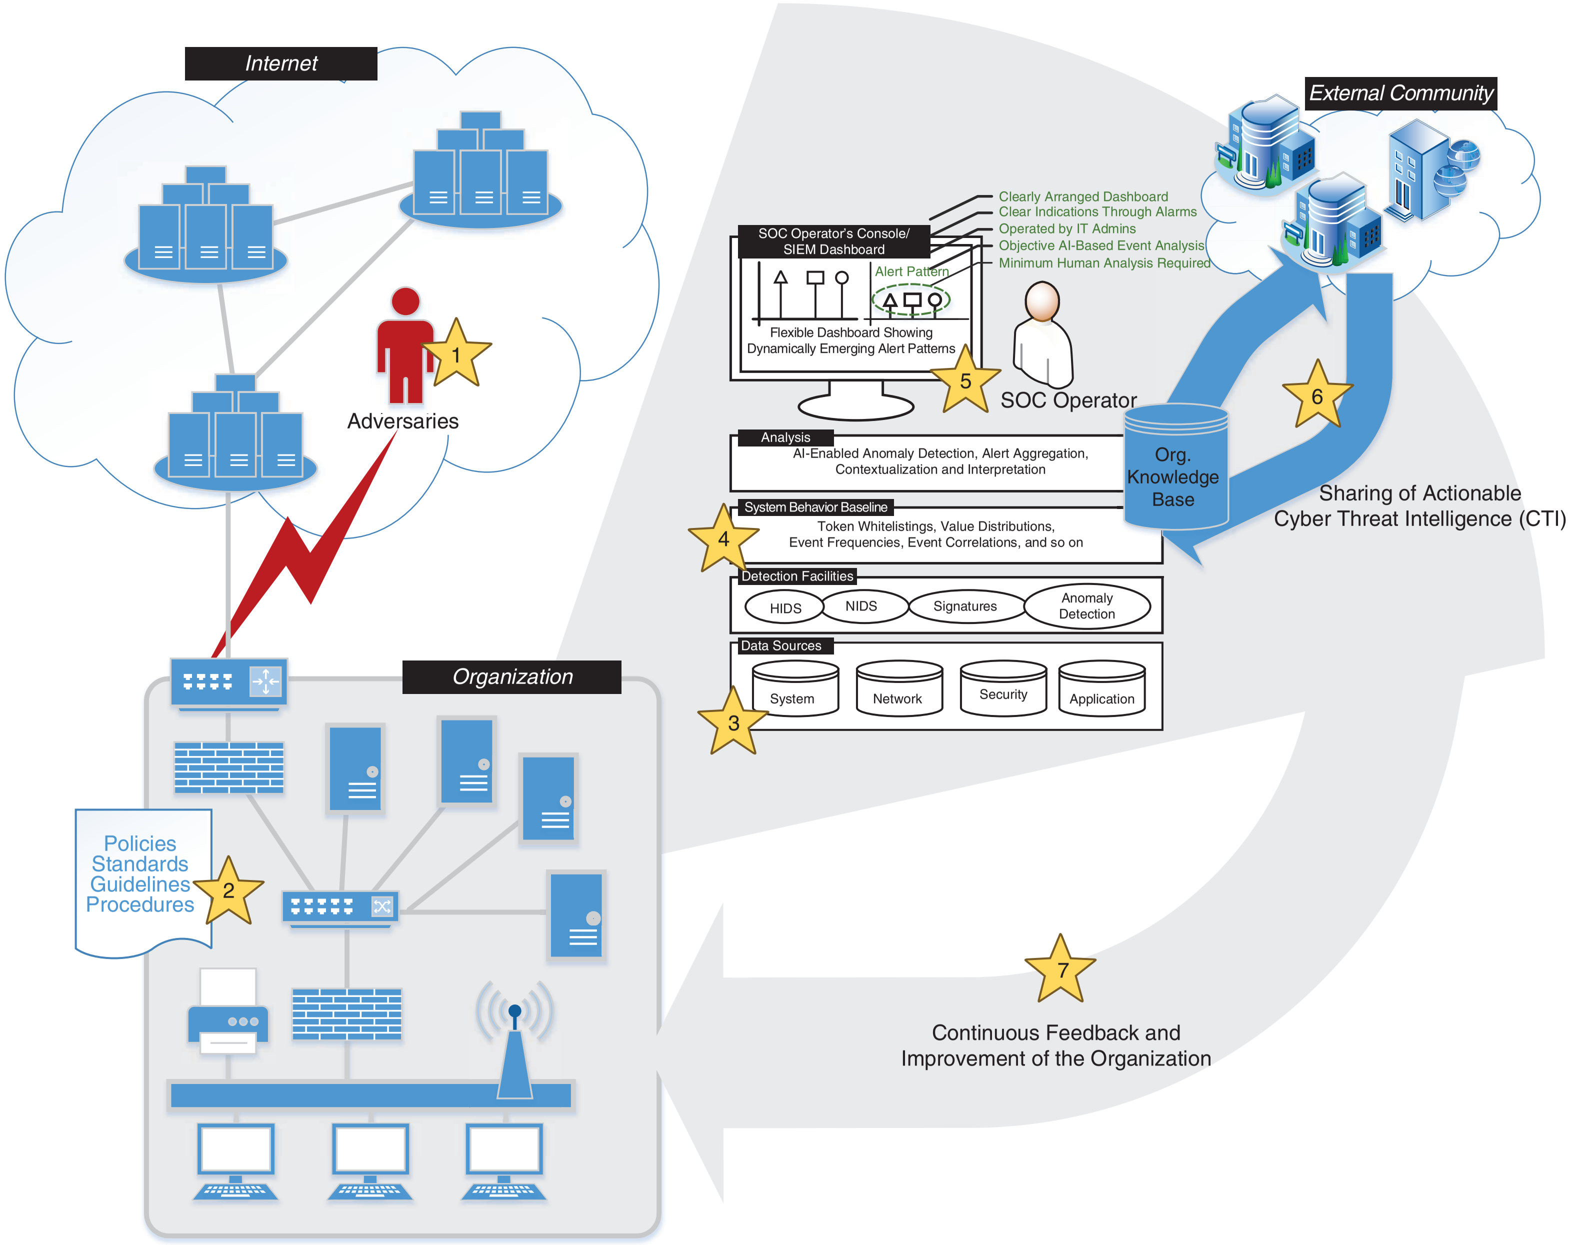
Internet
Adversaries
1
Organization
Policies
Standards
Guidelines
Procedures
2
SOC Operator's Console/
SIEM Dashboard
Alert Pattern
Flexible Dashboard Showing
Dynamically Emerging Alert Patterns
5
Clearly Arranged Dashboard
Clear Indications Through Alarms
Operated by IT Admins
Objective AI-Based Event Analysis
Minimum Human Analysis Required
SOC Operator
External Community
6
Org.
Knowledge
Base
Sharing of Actionable
Cyber Threat Intelligence (CTI)
Analysis
AI-Enabled Anomaly Detection, Alert Aggregation,
Contextualization and Interpretation
System Behavior Baseline
Token Whitelistings, Value Distributions,
Event Frequencies, Event Correlations, and so on
4
Detection Facilities
HIDS
NIDS
Signatures
Anomaly
Detection
Data Sources
System
Network
Security
Application
3
7
Continuous Feedback and
Improvement of the Organization

Figure 3. Intrusion detection and threat-hunting methodology in an enterprise environment and associated principles: 1) get to know your enemy, 2) get to know thyself, 3) be open to all sorts of data, 4) analyze smart, 5) make sense out of data, 6) sharing means caring, and 7) learn and prepare for the next wave.

important than the others. They all work together hand-in-hand toward the same goal. If carefully implemented and adjusted, they will allow us to elevate intrusion detection to an unprecedented level.

Future research challenges are manifold. While there exist several solutions with respect to most of the introduced principles, they are neither broadly accepted yet nor do they complement each other in a seamless manner. Moreover, their further development and continuous adaptation to new constraints is a top priority. New emerging paradigms, including the Internet of Things and the ongoing adoption of cloud computing, make systems even more complex, which will offer novel opportunities for exploitation. Furthermore, the increasing interoperability contribute to not only more open and dynamic businesses but at the same time also increases the attack surface and will lead to novel business models for cybercriminals, too.

Fortunately, novel anomaly-based IDS provide an elegant means to cope with this increased threat level. A convenient way to try out a well-developed but cost-effective system is the open source project AMiner, available at https://github.com/ait-aecid/logdata-anomaly-miner. ■

Acknowledgments

This work was supported in part by the Austrian FFG project DECEPT under grant 873980, and in part by the European Union H2020 project GUARD under grant 833456 and the EDIDP project PANDORA under grant SI2.835928.

References

1. D. Miller, S. Harris, A. Harper, S. VanDyke, and C. Blask, *Security Information and Event Management (SIEM) Implementation*. New York: McGraw-Hill, 2011.
2. B. E. Strom, A. Applebaum, D. P. Miller, K. C. Nickels, A. G. Pennington, and C. B. Thomas, "MITRE ATT&CK: Design and philosophy," The MITRE Corporation, McLean, VA, Tech. Rep., 2018.
3. M. Christiansen, "Bypassing malware defenses," *SANS Inst. InfoSec Reading Room*, pp. 1–39, 2010.
4. M. Shashanka, M. Y. Shen, and J. Wang, "User and entity behavior analytics for enterprise security," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2016, pp. 1867–1874. doi: 10.1109/BigData.2016.7840805.
5. R. M. Lee and R. T. Lee, "SANS 2018 threat hunting survey results, SANS analyst program," SANS Inst., Bethesda, MD, Sept. 2018. [Online]. Available: http://staging-resources.malwarebytes.com/files/2018/09/Survey_ThreatHunting-2018_Malwarebytes.pdf
6. V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surveys* (*CSUR*), vol. 41, no. 3, pp. 1–58, 2009. doi: 10.1145/1541880.1541882.
7. "Mandiant M-Trends 2021," FireEye Inc, M-EXT-RT-US-EN-000372-02, Milpitas, CA, Special Rep., 2021. [Online]. Available: https://content.fireeye.com/m-trends/rpt-m-trends-2021
8. "2020 threat hunting report," Crowdstrike Overwatch Team, Sunnyvale, CA, 2020. [Online]. Available: https://go.crowdstrike.com/crowdstrike-2020-overwatch-threat-hunting-report.html
9. "The value of threat intelligence," Ponemon Inst. LLC, Traverse City, MI, 2019. https://stratejm.com/wp-content/uploads/2019/08/2019_Ponemon_Institute-Value_of_Threat_Intelligence_Research_Report_from_Anomali.pdf
10. T. Yadav and A. M. Rao, "Technical aspects of cyber kill chain," in *Proc. Int. Symp. Security Comput. Commun.*, Aug. 2015, pp. 438–452.
11. I. Mokube and M. Adams, "Honeypots: Concepts, approaches, and challenges," in *Proc. 45th Annu. Southeast Regional Conf.*, Mar. 2007, pp. 321–326.
12. S. Khalid, T. Khalil, and S. Nasreen, "A survey of feature selection and feature extraction techniques in machine learning," in *Proc. Sci. Inf. Conf.*, Aug. 2014, pp. 372–378.
13. C. Wagner, A. Dulaunoy, G. Wagener, and A. Iklody, Oct. 2016, "MISP: The design and implementation of a collaborative threat intelligence sharing platform," in *Proc. 2016 ACM Workshop on Inf. Sharing Collab. Security*, pp. 49–56.
14. S. Barnum, *Standardizing Cyber Threat Intelligence Information with the Structured Threat Information Expression (STIX)*. Bedford, MA: Mitre Corp., 2012, vol. 11, pp. 1–22.
15. B. J. Wood and R. A. Duggan, "Red teaming of advanced information assurance concepts," in *Proc. DARPA Inf. Survivability Conf. Expo. DISCEX'00*, Jan. 2000, vol. 2, pp. 112–118.

Florian Skopik is head of the cybersecurity research program at the Austrian Institute of Technology, Wien, 1210, Austria. He coordinates national and large-scale international research projects, as well as the overall research direction of the team. His main interests are centered on critical infrastructure protection, intrusion detection, and national cybersecurity. Skopik received a Ph.D. in computer science from Vienna University of Technology. He is a member of various conference program committees (e.g., Association for Computing Machinery, Symposium on Applied Computing, International Conference on Availability, Reliability and Security, International Conference on Critical Information Infrastructure Security), editorial boards, and standardization groups, such as ETSI TC Cyber, IFIP TC11 WG1, and OASIS CTI. He frequently serves as a reviewer for numerous high-profile journals, including Elsevier's *Computers & Security* and *ACM Computing Surveys*. He is a Senior Member of IEEE. Contact him at florian.skopik@ait.ac.at.

Markus Wurzenberger is a scientist and project manager at the Austrian Institute of Technology, Wien, 1210, Austria. His main research interests are log data analysis with a focus on anomaly detection (AD) and cyberthreat intelligence. Wurzenberger received a Ph.D. in computer science from Vienna University of Technology. Besides the involvement in several national and international research projects, Wurzenberger is one of the key researchers working on the Austrian Institute of Technology's AD project Automatic Event Correlation for Incident Detection (AECID). Among the most prominent solutions developed within this project, Wurzenberger and his team created AMiner, a software component for log analysis, which implements several AD algorithms and is included as package in the official Debian distribution. Contact him at markus.wurzenberger@ait.ac.at.

Max Landauer is a scientist at the Austrian Institute of Technology, Wien, 1210, Austria. His main research interests are log data analysis, anomaly detection (AD), and cyberthreat intelligence. Landauer is a Ph.D. candidate in computer science at the Vienna University of Technology. His Ph.D. studies are a cooperative project between the Vienna University of Technology and the Austrian Institute of Technology. For his dissertation, he is working on an automatic threat intelligence mining approach that extracts actionable cyberthreat intelligence from raw log data. The goal of this research is to transform threat information shared by different organizations into abstract alert patterns that allow detection and classification of similar attacks. Landauer has multiple years of experience with nationally and internationally funded projects in numerous areas, including machine learning, artificial intelligence, cyberphysical systems, and digital service chains. Contact him at max.landauer@ait.ac.at.

Digital Object Identifier 10.1109/MSEC.2021.3103620